



TechBriefs

National Aeronautics and
Space Administration



Electronic Components and Circuits



Electronic Systems



Physical Sciences



Materials



Computer Programs



Mechanics



Machinery



Fabrication Technology



Mathematics and Information Sciences



Life Sciences

INTRODUCTION

Tech Briefs are short announcements of new technology derived from the research and development activities of the National Aeronautics and Space Administration. These Briefs emphasize information considered likely to be transferable across industrial, regional, or disciplinary lines and are issued to encourage commercial application.

Availability of NASA Tech Briefs and TSP's

Distribution of NASA Tech Briefs, a monthly periodical publication, is limited to engineers in U.S. Industry and to other domestic technology transfer agents. Requests for individual Tech Briefs or for Technical Support Packages (TSP's) announced herein should be addressed to

NASA Center for AeroSpace Information
Technology Transfer Office
800 Elkrige Landing Rd.
Linthicum Heights, MD 21090-2934
Telephone No. (301) 621-0245

Please reference the three-letter, five-digit control number located at the end of each Tech Brief. Information on NASA's Technology Utilization Program, its documents, and services is also available at the same facility.

Technology Utilization Officers and Patent Counsels are located at NASA field installations to provide technology-transfer access to industrial users. Inquiries can be made by writing to NASA field installations listed below.

Technology Utilization Officers and Patent Counsels

Ames Research Center
Technology Utilization Officer
Mail Code 223-3
Moffett Field, CA 94035

Patent Counsel
Mail Code 200-11
Moffett Field, CA 94035

Goddard Space Flight Center
Technology Utilization Officer
Mail Code 702-1
Greenbelt, MD 20771

Patent Counsel
Mail Code 204
Greenbelt, MD 20771

Lyndon B. Johnson Space Center
Technology Utilization Officer
Mail Code IC-4
Houston, TX 77058

Patent Counsel
Mail Code AL3
Houston, TX 77058

John F. Kennedy Space Center
Technology Utilization Officer
Mail Stop PT-PMO-A
Kennedy Space Center, FL 32899

Patent Counsel
Mail Code PT-PAT
Kennedy Space Center, FL 32899

Langley Research Center
Technology Utilization Officer
Mail Stop 143
Hampton, VA 23665

Patent Counsel
Mail Code 279
Hampton, VA 23665

Lewis Research Center
Technology Utilization Officer
Mail Stop 7-3
21000 Brookpark Road
Cleveland, OH 44135

Patent Counsel
Mail Code LE-LAW
21000 Brookpark Road
Cleveland, OH 44135

Jet Propulsion Laboratory
Technology Utilization Officer
Mail Stop 156-211
4800 Oak Grove Drive
Pasadena, CA 91109

NASA Resident Office-JPL
Technology Utilization Officer
Mail Stop 180-801
4800 Oak Grove Drive
Pasadena, CA 91109

Patent Counsel
Mail Code 180-801
4800 Oak Grove Drive
Pasadena, CA 91109

George C. Marshall Space Flight Center
Technology Utilization Officer
Code AT01
Marshall Space Flight Center,
AL 35812

Patent Counsel
Mail Code CC01
Marshall Space Flight Center,
AL 35812

John C. Stennis Space Center
Technology Utilization Officer
Code HA-30
Stennis Space Center, MS 39529

NASA Headquarters
Technology Utilization Officer
Code CU
Washington, DC 20546

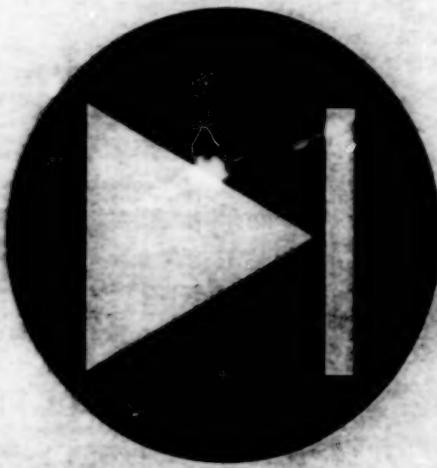
Assistant General
Counsel for Patent Matters
Code GP
Washington, DC 20546

Dryden Flight Research Center
Technology Utilization Officer
M/S D21-31
Bldg. 4832 Whse 7
Lilly Dr.
Edwards, CA 93523

National Aeronautics and
Space Administration5 **Electronic Components and Circuits**13 **Electronic Systems**21 **Physical Sciences**29 **Materials**39 **Computer Programs**47 **Mechanics**53 **Machinery**59 **Fabrication Technology**65 **Mathematics and Information Sciences**

This document was prepared under the sponsorship of the National Aeronautics and Space Administration. Neither the United States Government nor any person acting on behalf of the United States Government assumes any liability resulting from the use of the information contained in this document, or warrants that such use will be free from privately owned rights.

BLANK PAGE



Electronic Components and Circuits

Hardware, Techniques, and Processes

- 7 Split-Waveguide Mounts for Submillimeter-Wave Multipliers and Harmonic Mixers
- 8 Diffusion-Barrier Contacts for Solar Cells
- 8 CCD Detects Two Images in Quick Succession
- 9 Antennas for Receiving Signals Broadcast via Satellites
- 10 Algorithm for Detecting dc Series Arcs
- 10 Focal-Plane Analysis for Calculating Antenna Gain
- 11 Temperature-Compensated Sapphire Microwave Resonator

BLANK PAGE

Split-Waveguide Mounts for Submillimeter-Wave Multipliers and Harmonic Mixers

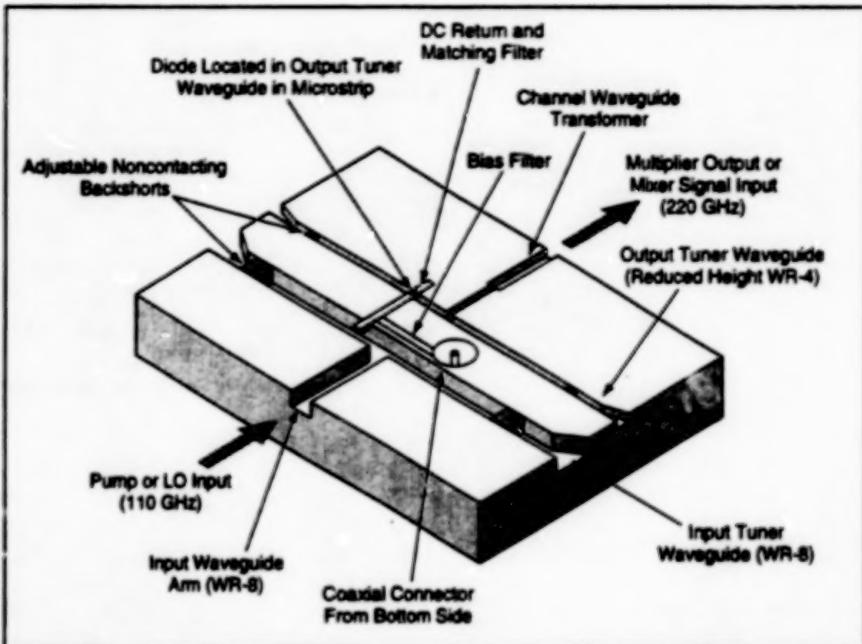
Advantages include ease of fabrication, reliability, and tunability.

NASA's Jet Propulsion Laboratory, Pasadena, California

A novel variation of a split-waveguide mount for millimeter- and submillimeter-wavelength frequency multipliers and harmonic mixers has been developed. The mount was designed to offer a wide range of available matching impedances, while maintaining a relatively simple fabrication sequence. The wide tuning range is achieved with separate series and parallel elements, which consist of two pairs of noncontacting sliding backshorts, at the fundamental and harmonic frequencies. The backshorts are arranged side by side between the input/output coupling arms, which makes for a very compact design with easy access to the tuner adjustments. Waveguide arms perpendicular to the tuner channels provide in-line input and output ports for multiplication or in-line signal and local oscillator (LO) ports for harmonic mixing. The mount is fabricated in two pieces, which are nearly mirror images of each other. The design is meant to be implemented with ordinary machining techniques — principally end-milling, sawing, or scribing the block halves. All six waveguide channels are in the same plane and no electroforming is required.

These waveguide mounts are particularly well-suited for circuits in which the electrically nonlinear elements are incorporated into stripline or suspended substrate circuitry rather than the more traditional, but circuit limiting, whisker-contact structures (although these can be accommodated as well). Both theoretical and experimental measurements of the mount indicate that the generic design enables matching over a wide range of impedances at both the fundamental and harmonic frequencies.

The figure shows the layout of a 110- to 220-GHz frequency doubler using the new mount design. The overall dimensions of the block (with the two nearly-mirror-image block halves assembled) are 29 × 14 × 19 mm, where 14 mm is the distance between the input and output waveguide flanges. The 110-GHz input pump signal enters via a full height WR-8 waveguide (1.02 × 2.03 mm). The two opposing input tuner arms, which are also WR-8 waveguides, contain the noncon-



The Block Containing Waveguide Channels and a few discrete components is one of two nearly-mirror-image blocks that are assembled to construct a frequency doubler.

tacting sliding backshorts, which provide series and parallel reactive tuning elements for matching at the input frequency. The backshorts are adjusted by use of micrometer heads. From the input tuner waveguide, the pump is coupled to a quartz microstrip probe, then along a microstrip harmonic separating filter to the multiplier diode, which is located across the output tuner waveguide. The filter is composed of chrome/gold low-pass elements on a fused substrate 0.152 mm thick and 0.33 mm wide. The distance between the input and output tuner waveguides is 2.50 mm. The filter extends beyond the output waveguide to allow additional electrical matching of both the input and output harmonic frequencies or possible idlers. In the doubler, the extended filter region provides an electrical short at the waveguide wall for both fundamental and second harmonic frequencies.

Output power leaves the block via an electric-field-plane waveguide arm perpendicular to the tuner arms. The output tuner waveguide, like the input, contains two noncontacting sliding backshorts, which provide series and parallel reactive matching elements at the output fre-

quency of 220 GHz. The output impedance matching range is broadened by using reduced, rather than full, height waveguide (0.28 × 1.09 mm).

A built-in channel-waveguide-transformer brings the output guide up to standard height WR-4 (0.54 × 1.09 mm). The distances of the input and output port electric-field-plane arms from the input coupling probe and the multiplier diode are approximately one half of a guide wavelength at the respective center frequencies of operation. An additional filter, between the input/output tuner guides and perpendicular to the harmonic separating filter provides biasing capability for the diode. This filter ends in a wire bond to the center pin of an SMA coaxial connector which extends up from the bottom of the block.

This work was performed by Antti Raisanen, Debaben Choudhury, Robert J. Dengler, John E. Oswald, and Peter H. Siegel of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

NPO-19248

Diffusion-Barrier Contacts for Solar Cells

The barriers enhance performance at low temperatures and low light intensities.

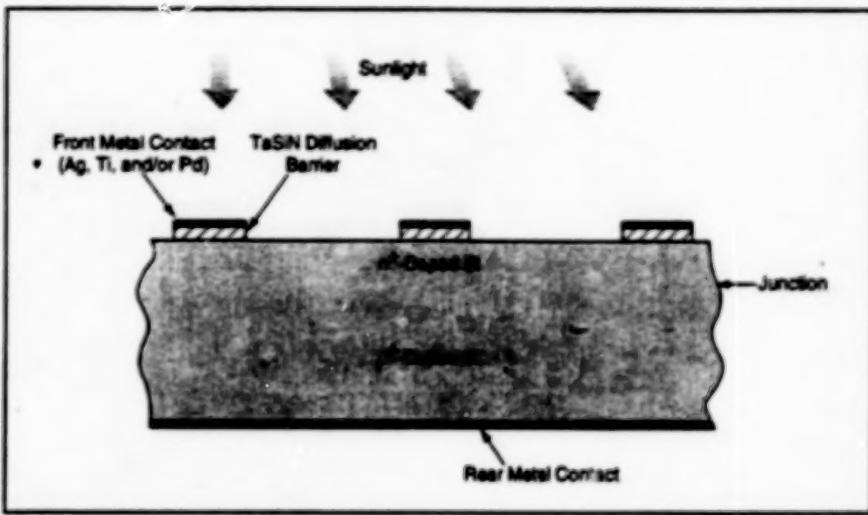
NASA's Jet Propulsion Laboratory,
Pasadena, California

Diffusion-barrier contacts can be incorporated into silicon-based solar photovoltaic cells to reduce degradation of performance at low temperatures and low light intensities. Heretofore, the low-temperature, low-light-intensity performances of such cells have been degraded by interactions that occur between metal contacts and silicon during fabrication.

The present diffusion-barrier contacts are thin, electrically conductive strips of TaSIN, which are deposited on the front surface of each affected cell. The strips of the front-surface metal contact grid are then deposited on the diffusion-barrier contacts (see figure). The diffusion-barrier contacts prevent diffusion of metal into the silicon at the high processing temperature.

This work was done by Paul M. Stella, Frederick S. Pool, and Marc Nicolet of Caltech and Peter A. Iles of Applied Solar Energy Corp. for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

NPO-19513

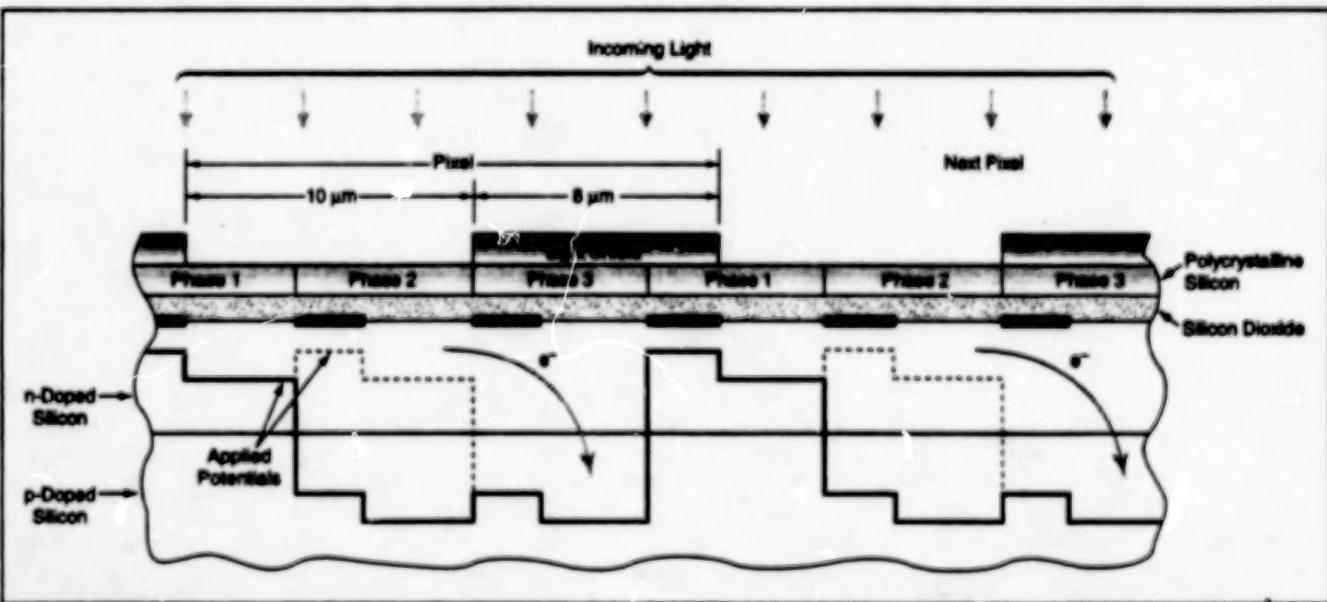


Electrically Conductive Diffusion Barriers of TaSIN prevent diffusion of metal from the overlying metal contacts into the underlying silicon during processing at high temperature, thereby improving performance during subsequent use in low-intensity light at low temperature.

CCD Detects Two Images in Quick Succession

Readout is performed slowly to minimize noise.

NASA's Jet Propulsion Laboratory,
Pasadena, California



This Modified Three-Phase CCD is operated in synchronism with two pulses of light to detect two images. The first image is formed in phases 1 and 2, then quickly transferred to phase 3. The second image is formed and stored in phases 1 and 2. Then the two images are read out in a slow scan.

A prototype special-purpose charge-coupled device (CCD) is designed to detect two $1,024 \times 1,024$ -pixel images in rapid succession (typically, separated by < 100 ns). When fully developed, this CCD would be operated in synchronism with a

pulsed laser, stroboscope, or other pulsed source of light to form pairs of images of rapidly moving objects.

The basic design of this CCD is similar to that of a standard 3-phase, buried-channel CCD, except that within each

pixel, an implant is added to each phase, and a light shield is added to phase 3 (see figure). A typical sequence of operations involves two flashes of light separated by the desired interval of the order of 100 ns. During the first flash, the first image is

formed in phases 1 and 2 of the pixels (but not in phase 3 because of the light shields). During the interval of darkness between the first and second flashes, this first image is transferred to phase 3, where it is stored for subsequent readout.

The second image is formed in phases 1 and 2 during the second flash of light and is

then stored in phases 1 and 2 for subsequent readout. Thereafter, the first and second images are read out from the CCD by clocking phases 1, 2, and 3 of the pixels in the appropriate sequence. The readout time is chosen to be much longer than the interval between flashes (readout is performed in a slow scan) to minimize readout noise. The

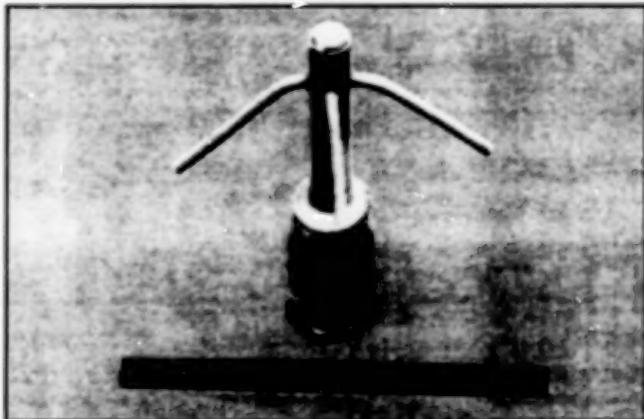
design level of readout noise is about 3 electrons, root mean square, per pixel.

This work was done by James R. Janesick and Andy Collins of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1]. NPO-19291

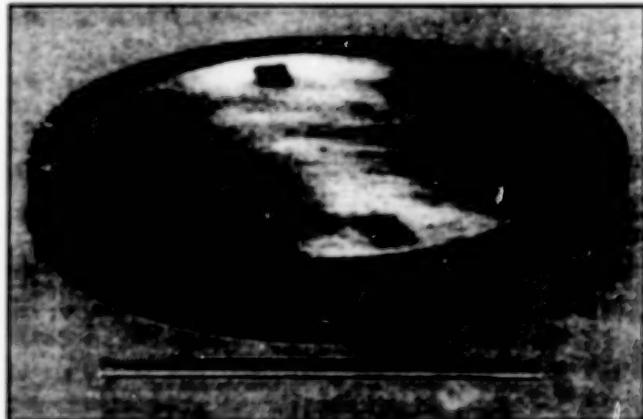
Antennas for Receiving Signals Broadcast via Satellites

Small, inexpensive units have been designed for outdoor and indoor uses.

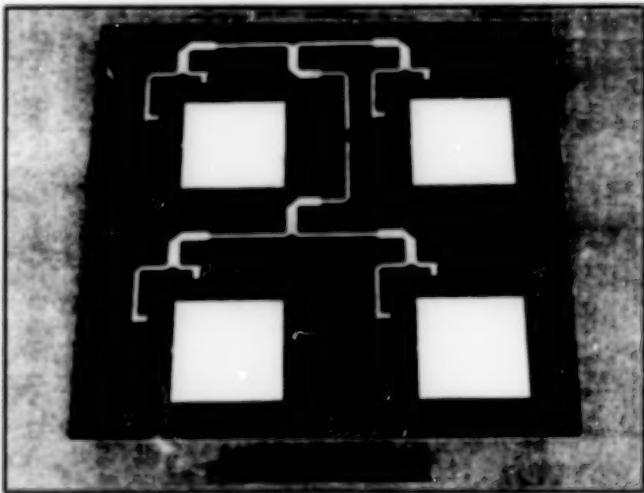
NASA's Jet Propulsion Laboratory, Pasadena, California



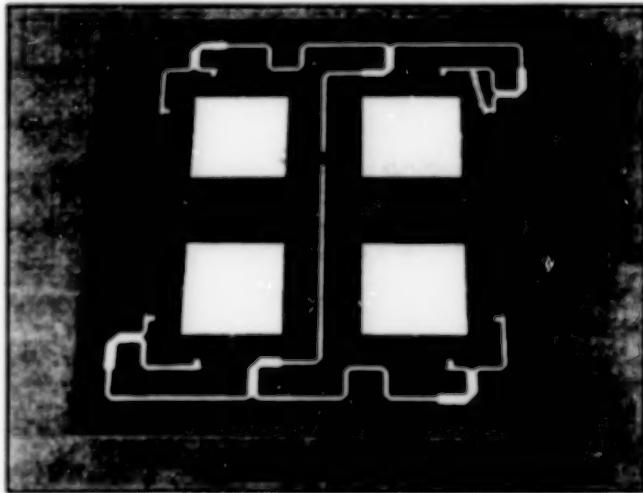
CROSSED DROOPING DIPOLES



TM₂₁-MODE CIRCULAR MICROSTRIP PATCH ANTENNA



SQUARE PATCHES WITH CONVENTIONAL FEED



SQUARE PATCHES WITH SEQUENTIAL FEED

These Small, Inexpensive Antennas are suitable for reception of satellite broadcasts in right-hand circular polarization at a frequency of 2.05 GHz.

Four different antennas have been designed for use in receiving radio AM and FM signals broadcast via Earth-orbiting satellites at a carrier frequency of 2.05 GHz in right-hand circular polarization. These antennas are small, lightweight, inexpensive units with low-to-medium-gain, quasi-omnidirectional radiation patterns.

The first antenna, shown in the top left part of the figure, includes two crossed drooping dipoles protruding from a shaft, plus an impedance-matching tuning ring on the shaft. Each arm in one of the dipoles

is 1.2 in. (30.5 mm) long; each arm in the other dipole is 1.8 in. (45.7 mm) long. The height of the dipoles above the automobile roof or other ground plane can be adjusted to optimize performance. In one test, for example, the dipoles were set at a height of 3.2 in. (81.3 mm), yielding a radiation pattern that peaked at 60° below the zenith, with a peak gain of 4.8 dBic (where "dBic" denotes decibels with respect to isotropic, circularly polarized radiation).

The second antenna, shown in the top right part of the figure, includes a circular

conductive TM₂₁-mode-radiating patch of 5.2-in. (132-mm) diameter on a honeycomb dielectric plate of 6-in. (152-mm) diameter and 0.5-in. (12.7-mm) thickness. The patch is connected, at four locations near its periphery, to four ports of a printed-circuit-board stripline feed network that provides the amplitude and phase relationships for reception of right-hand circular polarization. This antenna is also intended for outdoor use. In a test, this antenna provided a peak gain of 6.5 dBic at an angle of 36° below the zenith.

The other two antennas are medium-gain units designed for indoor use. These antennas contain 2×2 arrays of printed-circuit square conductive patches and microstrip feed lines that provide the required amplitude and phase relationships. Both of these antennas are fabricated on dielectric substrates of 0.126-in. (3.2-mm) thickness. These antennas differ from each other in their microstrip feed configurations and in their lateral dimensions, which are 8.5×8.5 in. (216 \times 216 mm) for one and 8×8 in. (203 \times 203 mm) for the other. The measured broadside

gains of the units were 12.0 and 13.7 dBic, respectively.

This work was done by Te-Kao Wu and John Huang of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1]. NPO-19362

Algorithm for Detecting dc Series Arcs

Load currents are analyzed to identify fluctuations characteristic of arcs.

An algorithm has been developed for use in automatic detection of dc series arcs in power cables, cable terminations, and junction boxes. (The algorithm can also sometimes detect arcs inside electrical loads.) A system in which the algorithm can be utilized is one that receives its power from a microprocessor-based remote power controller that includes a load-current sensor. The algorithm analyzes the digitized load-current readings for fluctuations characteristic of dc series arcs.

An algorithm, as distinguished from something simpler, is needed for automatic detection of dc series arcs because the problem of identifying the arcs from load-current readings is not as simple as one might initially be inclined to think. A dc series arc is an arc that occurs whenever a switch carrying dc opens or a wire carrying dc ruptures. The arc continues until the arc current ceases to flow. Unlike an arc to ground (which is a short circuit and therefore easily detectable by conventional overcurrent-sensing circuitry), a dc series arc does not give rise to a large overcurrent;

indeed, the current in a dc series arc can closely resemble normal current.

Experiments have shown that a dc series arc is characterized by a step (an increase or decrease, depending on the type of equipment) of typically about 8 percent in the load-current readings at the inception of the arc, followed by an increase in the noise component of the load current. The algorithm is based on this characteristic, and it works as follows: During the first few minutes of presumably normal operation, the digitized load-current readings are monitored to determine an initial background noise level. Thereafter, the monitoring continues, the background noise level is updated from time to time in coordination with the operation of the power-controlling microprocessor, and the noise level measured during each 10-ms interval is compared with the background level. A dc series arc is tentatively deemed to have begun whenever this short-term noise rises to several thousand times the background level during one 10-ms interval.

Marshall Space Flight Center, Alabama

Once an arc has thus been tentatively identified and for the next 300 ms, the number of 10-ms intervals during which this short-term noise exceeds 3 times the background noise level is counted. If the count exceeds 9, then the tentative identification is deemed to be confirmed.

Hazards that can be prevented by timely detection of dc series arcs include generation of smoke and poisonous gases, and ignition of nearby combustible materials. The algorithm has been optimized to reduce the false-alarm rate to about once per century. Of course, the algorithm offers no protection during the initial background-noise-monitoring interval.

This work was done by Greg L. Moore of Micon Engineering, Inc., for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-26340.

Focal-Plane Analysis for Calculating Antenna Gain

Computation time is reduced substantially below that of the traditional method.

An improved method has been devised for calculating the gain of an antenna system that comprises a main reflector and an array of feed elements. In comparison with the traditional method, the improved method involves much less computation; this is an important advantage when it is necessary to compute gains repeatedly, with slight variations in design at each iteration, in an effort to find an optimum design.

In the traditional method, one calculates the gain from a perspective of transmission. One begins by calculating the current density produced on the reflector surface by each element of the array, then uses the classic radiation surface integral

with an integrand proportional to $[j \text{current density}] \times [\text{a direction-dependent phase factor}]$ to obtain the far electromagnetic field generated by the element. The gain is then calculated from the sum of the far fields generated by all of the elements. The calculation of far fields for all of the elements can consume much time; the amount of computation time can become prohibitive when it is necessary to perform the far-field calculations for evaluating a number of alternative geometric configurations for the array.

In the improved method, one calculates the gain from the perspective of reception, and it is not necessary to repeat the far-

NASA's Jet Propulsion Laboratory, Pasadena, California

field calculation for each different geometric configuration of the array configuration. Instead, one computes a near reflected field that corresponds to a far field incident from a specified direction of reception, and it is necessary to perform this part of the calculation only once for each combination of reflector-surface configuration and receiving direction; it is not necessary to repeat this part of the calculation as long as the reflector-surface configuration and far-field direction remain the same.

The improved method is based on the Lorentz reciprocity theorem and, more particularly, on one of its consequences; namely, that the antenna gain in transmis-

sion or reception (they are the same) can be computed from a convolution of two fields in the focal plane of the reflector. One of these fields is the near field mentioned above; more specifically, the focal-plane field produced by a plane wave

incident on the reflector from the specified direction of reception. The other field is the focal-plane field produced by the array of feed elements excited with currents in the specified amplitudes and phases.

This work was done by Paul W. Cramer, William A. Imbriale, and Sembiam R. Rengarajan of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1]. NPO-19460

Temperature-Compensated Sapphire Microwave Resonator

Ultrastable operation can be achieved at relatively low cost.

A sapphire-dielectric-ring microwave resonator that operates in a "whispering-gallery" electromagnetic mode features a differential-thermal-expansion design that provides temperature compensation for ultrahigh frequency stability. Heretofore, sapphire-ring resonators have been temperature-compensated by addition of paramagnetic impurities, but this approach has proven effective only at temperatures ≤ 6 K. The present resonator is a prototype of a type of sapphire resonator designed to minimize frequency fluctuations caused by temperature fluctuations at a nominal temperature equal to or even somewhat greater than the temperature of liquid nitrogen (77 K). The ancillary equipment needed for operation in this higher temperature range is smaller and less expensive, and liquid nitrogen can be used as the coolant.

This resonator includes an upper and a lower sapphire ring separated by a small gap and spaced apart by a central copper post that is almost as tall as the rings are. The use of the small gap preserves the high Q (quality factor, a measure of the sharpness of resonance, defined as $2\pi \times$ the energy stored in the electromagnetic field in the resonator + the energy dissipated per cycle of oscillation) of the whispering-gallery modes of interest while giving a steep variation of frequency with the size of the gap. The tall copper post gives a relatively large motion due to thermal expansion. The central position of the copper post minimizes its tendency to degrade Q because the electromagnetic energy of the modes of interest is concentrated near the radially outermost surfaces of the sapphire rings.

Compensation occurs as follows: As the temperature increases, the mode frequencies decrease because of concomitant thermal expansion of the sapphire and a concomitant increase in the relative permittivity of the sapphire. However, because the thermal expansion of the copper post exceeds that of the sapphire, the gap widens. The resulting increase in the volume of the gap (relative permittivity of vacuum = 1, as compared with relative

NASA's Jet Propulsion Laboratory, Pasadena, California

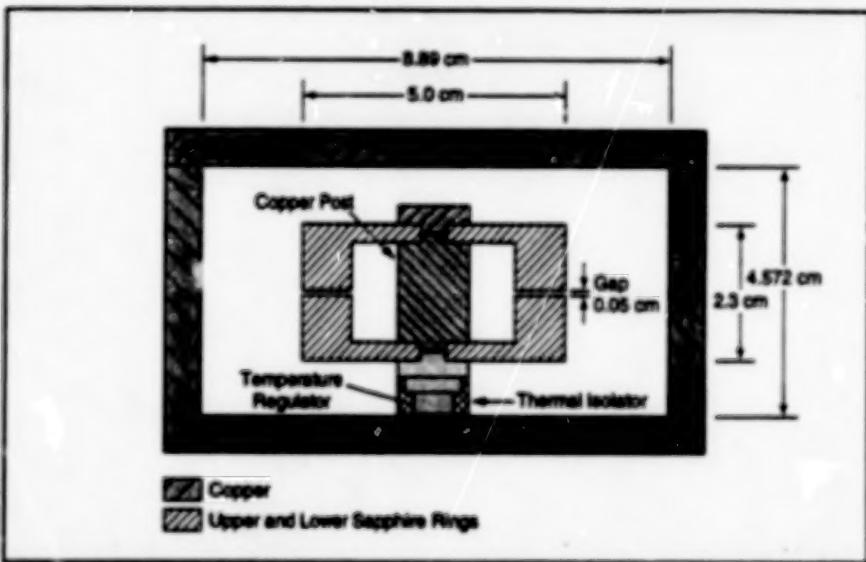


Figure 1. The Thermal Expansion of the Copper Post causes the frequencies of the electromagnetic modes of the resonator to increase with temperature, while thermal effects in the sapphire cause the frequencies to decrease with temperature. These opposing effects can be utilized for temperature compensation.

permittivity of sapphire = 10) tends to increase the frequencies of the modes. At some temperatures, these opposing effects can cancel each other; in other words, result in complete temperature compensation.

Figure 2 shows the measured temperature dependence of frequency for a mode with a nominal frequency of 7.23 GHz that is compensated at a temperature of 87.09 K. In this case, one could achieve a frequency stability characterized by a change in fractional frequency no greater than 10^{-14} during a typical observation interval 100 s long if one could maintain the temperature stable to within 43 μ K. This requirement can be satisfied by use of conventional temperature-regulation equipment developed for frequency-standard oscillators. The projected fractional frequency instability of 10^{-14} of this resonator is less than that of a typical highest-quality quartz resonator, which is about 2×10^{-13} .

This work was done by G. John Dick and David G. Santiago of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

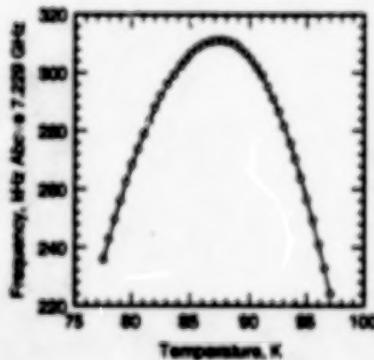


Figure 2. Complete Temperature Compensation of a 7.23-GHz mode of this resonator can be achieved by operating at 87.09 K, where the slope of the frequency-vs.-temperature curve goes to zero.

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, NASA Resident Office-JPL [see page 1]. Refer to NPO-19414.

BLANK PAGE



Electronic Systems

Hardware, Techniques, and Processes

- 15 Automated Monitoring of Dielectric Properties of Tree Trunks
- 16 Mapping Pixel Windows to Vectors for Parallel Processing
- 17 Coherent-Phase Monitoring of Cavitation in Turbomachines
- 17 Modular, Hierarchical Learning by Artificial Neural Networks
- 18 Ground-Based Calibration of a Microwave Landing System

BLANK PAGE

Automated Monitoring of Dielectric Properties of Tree Trunks

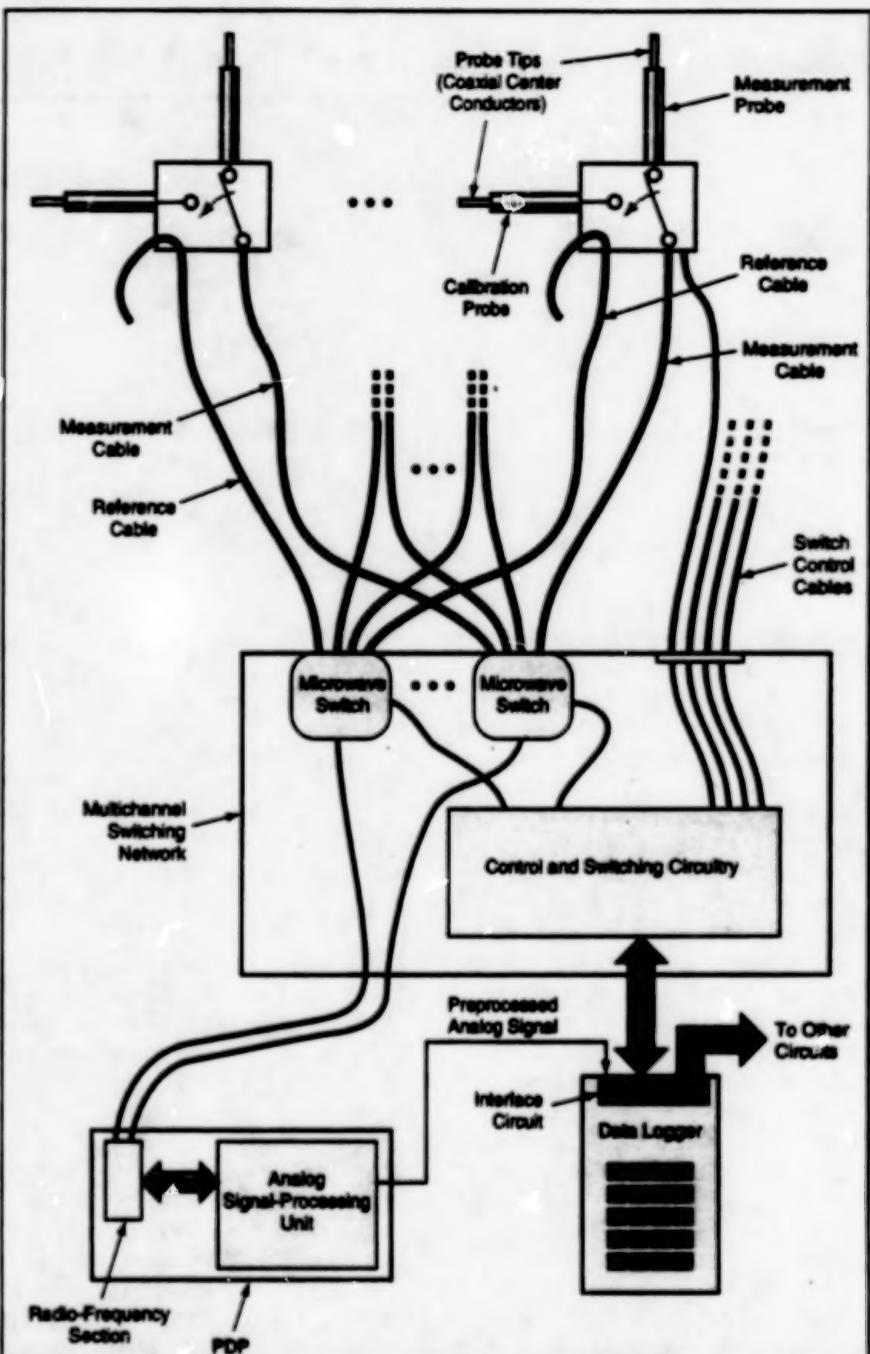
A semiautomated instrumentation system can be left unattended to monitor several trees.

NASA's Jet Propulsion Laboratory,
Pasadena, California

A semiautomated instrumentation system called a "dielectric monitoring system" (DMS) has been developed for measuring the microwave permittivities of selected components of plants — in particular, of active xylems in tree trunks. The system is built around a modified version of a commercial instrument, called a "portable dielectric probe" (PDP), which includes a coaxial probe tip that is inserted into the xylem or other component to be monitored. The DMS includes eight measurement channels for nearly simultaneous monitoring of as many as eight trees. After being set up in the field, the DMS can be left unattended to record the measurement data; thus, the DMS makes it possible to perform continuous monitoring of time-varying dielectric properties of trees in remote forest regions, where previously, the labor-intensive nature of the measurements made such continuous monitoring impractical.

The design of the DMS incorporates the PDP into a complete monitoring system. The PDP operates in conjunction with a multichannel switching unit and a data logger (see figure). The PDP includes a radio-frequency section and an analog signal-processing unit that provide both the signals needed at the probe tips for measuring permittivities and a preprocessed analog signal for the data logger. The multichannel switching unit provides the multichannel monitoring capability and provides for automatic monitoring of a dielectric calibration reference (usually air) for each of the measurement channels. The data logger includes an interface circuit that automatically controls both the channel switching and a data-storage medium that can hold data from as many as 120,000 measurements. The data logger also includes a 12-bit analog-to-digital converter for high resolution of stored measurement data.

The center conductor of the coaxial probe tip is shown extended, whereas the center conductor of the coaxial probe tip of the unmodified PDP was flush with the flat end of the probe cylinder. The extension of the center conductor increases the reliability of contact between the probe and the specimen; it also increases the measurement volume, thereby providing greater measurement sensitivity plus assurance that the measurement samples the dielectric properties throughout the



This Instrumentation System is set up with coaxial probes inserted in tree trunks to measure dielectric properties. The system can be left to operate unattended to gather data on permittivities as a function of time.

active xylem or other layer to be probed.

Each measurement channel consists of a switch control cable and a pair of microwave cables, one of which is connected to a two-position microwave switch. The microwave switch is toggled between each of two coaxial probe tips. For each measurement channel, one

measurement probe is mounted in the medium of interest, e.g., a tree trunk, while the other calibration probe is left exposed to the air to provide measurement of a reference material of known permittivity. The switch control cable allows for control of the microwave switch so that observations may be

obtained from each of these two probe tips. The reference coaxial cable is the same electrical length as the measurement cable-switch-probe tip ensemble. This is required for the PDP to properly

process the RF signal and determine the reflection coefficient observed at the probe tips.

This work was done by Kyle C. McDonald and William Chun of Caltech for

NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].
NPO-19338

Mapping Pixel Windows to Vectors for Parallel Processing

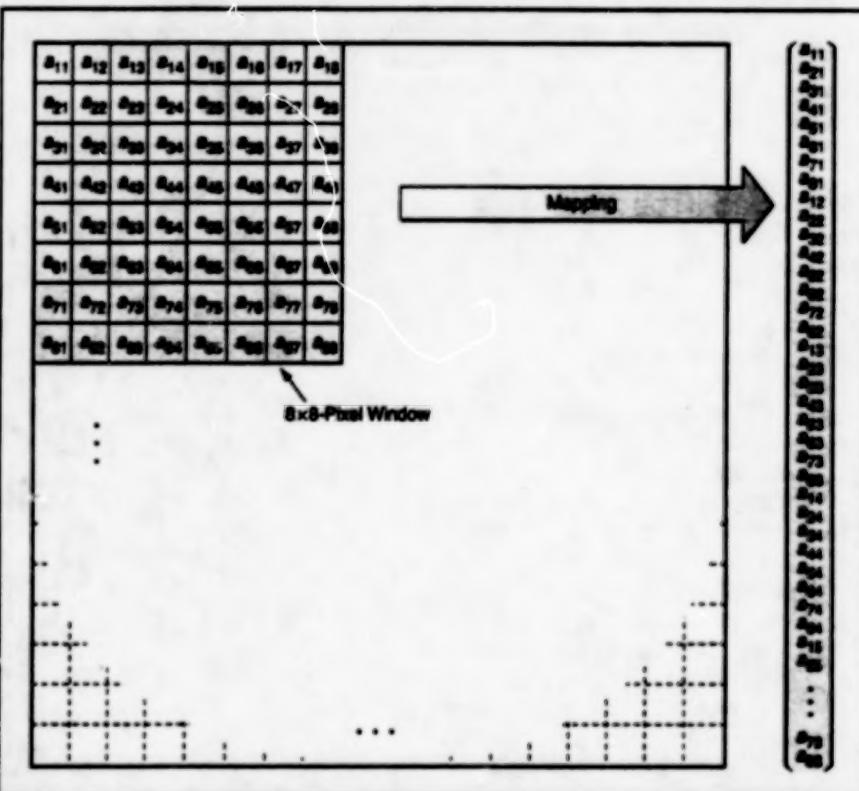
The mapping is performed by matrices of transistor switches.

Arrays of transistor switches have been devised for use in forming simultaneous connections from (1) a square subarray (window) of $n \times n$ pixels within an electronic imaging device that contains an $np \times np$ array of pixels to (2) a linear array of n^2 input terminals of an electronic neural network or other parallel-processing circuit. In mathematical terms, the transistor switches implement binary switching matrices that map the pixel window to an n^2 -dimensional vector (see figure); that is, the brightness-level output signal of each pixel constitutes one component of the vector, while the vector constitutes the collection of inputs to the parallel-processing circuit.

This mapping method helps to realize the potential for rapidity in parallel processing for such applications as enhancement of images and recognition of patterns. In providing simultaneous connections, this method overcomes the timing bottleneck of older multiplexing, serial-switching, and sample-and-hold methods. In addition, these older methods lack flexibility of connection patterns, whereas in the present method, the connection patterns can be varied at will by changing the matrix elements. Changing connection patterns can result in rearrangement of the sequence of components of the vector. Changing connection patterns can be used for a variety of purposes; for example, rotating or translating the window image fed to the parallel processor.

The mapping is performed in two stages. In the first stage, the mapping is from the pixels of the selected window to an intermediate window of $n \times n$ array of nodes. For this purpose, there must be a total of p^2 switches fanning in to each node from the corresponding pixels of the p^2 windows, and for each node, one of these switches is turned on to receive the signal from the corresponding pixel of the selected window. Thus, this mapping requires n^2p^2 switches and is represented by an $np \times np$ matrix, which is denoted as T .

NASA's Jet Propulsion Laboratory, Pasadena, California



All The Pixels in the Selected 8×8 Window are mapped simultaneously to positions along a 64-element linear array or, equivalently, to components of a 64-dimensional vector.

It is possible to select an $n \times n$ -pixel window that straddles one or more boundaries between adjacent $n \times n$ subdivisions of the overall $np \times np$ array of pixels. In such a case, the T -mapping yields vector components in the incorrect sequence. The second-stage mapping corrects the sequence. This mapping is represented by an $n^2 \times n^2$ matrix, M , and requires n^4 switches — one switch from each of the n^2 nodes to each of the n^2 input terminals of the parallel processor. M is also used to rearrange the components of the vector as needed for the desired transformation (e.g., rotation or translation) of the window image.

The "brute-force" method for effecting the same mapping with a flexible connection pattern requires $O(n^4p^4)$ switching

transistors [where $O(x)$ denotes a number proportional to a number of the order of x]. For reasonable values of n and p (e.g., $n, p \geq 8$), the circuitry needed to implement so many switches can be very complex and large, and its parasitic capacitance can be excessive. The present method can be implemented with simpler, smaller, and less expensive circuitry because it requires a smaller number of switching transistors; namely, $O(n^2p^2+n^4)$.

This work was done by Tuan A. Duong of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

NPO-19431

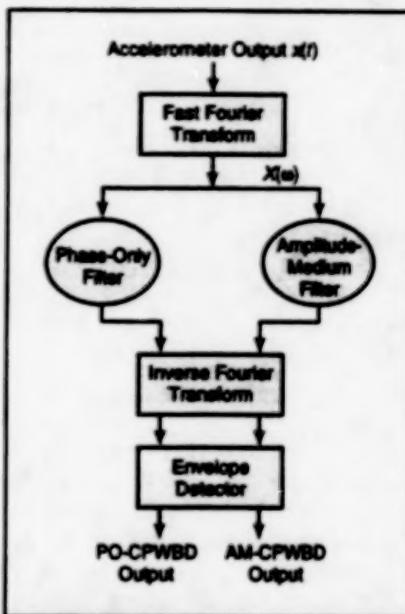
Coherent-Phase Monitoring of Cavitation in Turbomachines

Vibration signals are processed to extract measures of cavitation.

A digital electronic signal-processing system analyzes the outputs of accelerometers mounted on a turbomachine to detect vibrations characteristic of cavitation. It is desirable to detect and prevent cavitation in turbomachinery because cavitation can degrade performance, erode turbine blades, cause excessive structural vibrations, and eventually result in breakage.

The operation of this and older electronic cavitation-detection systems is based on a unique vibrational phenomenon associated with cavitation: when cavitation occurs in a rotary machine (e.g., a turbomachine), the wide-band noise generated by the collapse of cavitation bubbles becomes amplitude-modulated by the periodic components of motion associated with rotation (such as components at the fundamental frequency of rotation and its harmonics). This phenomenon gives rise to a periodicity in the accelerometer outputs; however, this periodicity is hidden in the sense that it is undetectable by ordinary linear spectral analysis.

Older cavitation-detection systems implement a technique called "full-wave rectification wide-band demodulation" (FWRWBD), which involves envelope detection of an isolated high-frequency band of an accelerometer output that includes a high-frequency, wide-band noise floor. The onset and existence of cavitation are determined by the recovery of periodicity in the detected signal, while the magnitude of the modulation is used to determine the severity of cavitation. Unfortunately, any discrete components (frequency components of structural vibrations) present in the high-frequency band in question give rise to peaks in the resulting demodulated signal. This limits the utility of FWRWBD in that it is difficult to distinguish between these peaks and



The CPWBD Algorithm recovers hidden periodicity, indicative of cavitation, from a noisy wide-band accelerometer output.

signal peaks contributed by recovered hidden periodicity.

The present system is designed to overcome the limitation imposed by interference from discrete components. It is based on the discovery that cavitation can be identified solely by its hidden coherent-phase information. The system digitally implements a technique called "coherent-phase wide-band demodulation" (CPWBD), using phase-only (PO) filtering along with envelope detection to search for the unique coherent-phase relationship associated with cavitation and to minimize the influence of large-amplitude discrete components.

The figure schematically illustrates the flow of information in the CPWBD algorithm. PO filtering is performed prior to a process in which a power spectral density is computed in the form of (a) a discrete

Marshall Space Flight Center,
Alabama

Fourier transform that is the product of (b) a frequency-band-limiting rectangular function and (c) the convolution of the discrete Fourier transform of the raw input signal with itself. The phase-only filter retains the coherent-phase relationship associated with the hidden periodicity that one seeks to identify. However, because of the amplitude-normalization effect of the phase-only filter, the system is no longer overwhelmed by independent discrete interference.

The amplitude of the hidden periodicity recovered via the PO part of the CPWBD algorithm is also subjected to some normalization. Therefore, the output of the PO part of the CPWBD algorithm should be viewed as a normalized "coherence" function for detecting the onset and existence of the wide-band-modulation (WBM) phenomena associated with cavitation. In order to detect the magnitude of the WBM signal and thereby gauge the intensity of cavitation, an "amplitude-medium" (AM) filter is introduced. This filter avoids the large vector-representation contribution of discrete interference while retaining the amplitude information about the degree of contribution toward the vector construction among various wide-band noise components. Therefore, the output of the AM part of the CPWBD algorithm should be viewed as a nonnormalized "spectrum" function for detecting the strength of the WBM signal associated with cavitation.

This work was done by Jen-Yi Jong of Al Signal Research, Inc., for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-26313.

Modular, Hierarchical Learning by Artificial Neural Networks

Learning is streamlined by imitating some aspects of biological neural networks.

A modular and hierarchical approach to supervised learning by artificial neural networks is a subject of continuing research. This approach leads to neural networks that are more structured than are neural networks in which all neurons are fully interconnected. These networks

utilize a general feedforward flow of information and sparse recurrent connections to achieve dynamical effects. The modular organization, the sparsity of modular units and connections, and the fact that learning is much more circumscribed than in networks of fully

NASA's Jet Propulsion Laboratory,
Pasadena, California

interconnected neurons are all attractive features for designing neural-network hardware.

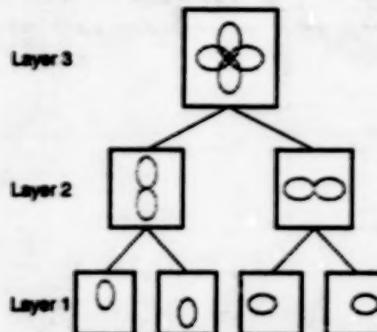
Previous approaches to learning by artificial neural networks have involved, variously, the concepts of back-propagation of error signals and gradient

descent of error measures in synaptic-connection-weight space. Although very powerful when applied to relatively simple problems, learning schemes based on these concepts alone break down as soon as sufficiently complex problems are considered. The modular, hierarchical approach was conceived to overcome this limitation.

This approach involves imitation of some aspects of learning schemes apparently followed by biological neural networks. Evolution seems to have overcome the obstacles inherent in gradient-descent learning. Learning in a biological neural network is never started from a tabula rasa. Rather, a high degree of structure is already present in the neural circuitry of a newborn organism.

Though little is known of the interaction between the prewired structure and the actual learning, it has been hypothesized that (1) complex learning tasks are broken up into simpler task modules that correspond to physical modules within a neural network, (2) learning, perhaps in different forms, can operate both within and across the gaps between modules, and (3) the modules can be organized in a hierarchical way. These hypotheses inspire the present modular, hierarchical approach.

The modular, hierarchical approach has been demonstrated and tested by applying it to the problem of making a neural network learn planar trajectories like circles and figure eights. The modules in each higher level of the hierarchy control the outputs of the modules one level below. The modules at the lowest level of the hierarchy contain small numbers of relatively simple neural units that



The Seven Modules are connected sparsely in a hierarchical network designed to learn a double-figure-eight trajectory.

can be regarded collectively, in some sense, as an oscillator that learns some simpler component of the overall motion (e.g., a sinusoid, a circle, or a Fourier component). Each lowest-level module could be a simple oscillator ring with two or three neurons, an odd number of inhibitory connections, and sufficiently high gains.

The figure illustrates an example in which these concepts are applied to learning a trajectory that comprises two perpendicular figure eights. In this case, each of four lowest-level modules contains an oscillator ring that learns one of the four loops of the trajectory. The parameters of the neural components in the modules can be adjusted by gradient descent to make the loops match those of the desired trajectory.

The second level of the hierarchy in this example contains two control modules. Each of these modules controls a distinct pair of oscillator modules from

the lowest level, so that it ends up producing a simple figure eight. Like the modules in the lower level, these modules can be oscillator rings and their parameters can be adjusted. In particular, after the learning process is completed, they should be operating in their high-gain regimes and the period of oscillation of each must equal the sum of the periods of the modules it controls.

The highest level of the hierarchy consists of another oscillatory and adjustable module that controls the two modules in the second level, so as to produce the double figure eight. This module must also end up operating in its high-gain regime. In general, the final output trajectory is a limit cycle because it is obtained by superposition of limit cycles of the various lower-level modules. If the various oscillators relax to their limit cycles independently of each other, it is essential to provide for adjustable delays between the various modules in order to get the proper harmony among the various phases.

In this way, a sparse network with 20 units or so can be constructed and made to execute a double figure eight.

This work was done by Pierre F. Baldi and Nikzad Toomarian of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, NASA Resident Office-JPL [see page 1]. Refer to NPO-19077.

Ground-Based Calibration of a Microwave Landing System

Reference data are acquired simultaneously from the Global Positioning System.

A system of microwave instrumentation and data-processing equipment has been developed to enable ground-based calibration of a microwave scanning-beam landing system (MSBLS) at distances of about 500 to 1,000 ft (150 to 300 m) from the MSBLS transmitting antenna. The system is needed to ensure the accuracy of the MSBLS near the touchdown point, without having to resort to the expense and complex logistics of aircraft-based testing. In the original application for which this system was developed, the MSBLS to be calibrated is the one that gives primary landing guidance to the space shuttle

orbiter during the final 20 miles (32 km) of approach. Modified versions of this system might prove useful in calibrating aircraft instrument landing systems.

As in the case of an aircraft-based system for calibrating the MSBLS at greater distances, the basic principle of this system is to use position data from the Global Positioning System (GPS) as references, with which the MSBLS position data are compared. However, instead of an MSBLS and a GPS receiver aboard an aircraft, this system includes an MSBLS and a GPS receiver installed in a trailer. The MSBLS and GPS receiving antennas are collocated

John F. Kennedy Space Center, Florida

at the top of a 50-ft (15-m) tower that is erectable from the trailer. Thus, the position of the collocated antennas simulates the position of a landing spacecraft or aircraft. To ensure the necessary precision of the GPS data, the system includes a GPS reference station located at a benchmark near the MSBLS transmitting antenna.

Once set up at the benchmark, the GPS reference station continues to receive signals from the GPS satellites, recording GPS code carrier-phase data throughout the calibration procedure. During the procedure, the MSBLS and GPS data received via the collocated antennas are recorded as the

tower is positioned with a towing vehicle on and near the runway, at various distances and directions from the MSEBLs transmitting antenna. Because the rates of acquisition of MSEBLs and GPS data are different, a 1-pulse-per-second time reference signal transmitted via the GPS is used to synchronize the collection of both sets of data. After an initialization interval of 10 min, the mobile

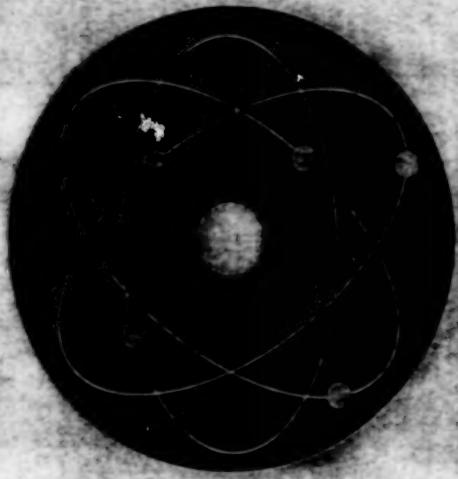
equipment can be moved to various locations, and data can be collected while the moves are underway. The GPS receivers provide real-time carrier phase position for the reference data, thus allowing the landing-system analysis to take place while acquiring the data.

This work was done by John J. Grizzas, Marshall M. Scott, Jr., and Alfred D. Willis

of Kennedy Space Center and Temel Erdogan and Rolando Reyes of I-NET, Inc. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Kennedy Space Center [see page 1]. Refer to KSC-11802.

BLANK PAGE



Physical Sciences

Hardware, Techniques, and Processes

- 23 Integrated Oxygen-Recovery System
- 24 Selective Oxidizer for Removal of Carbon Monoxide
- 24 Silicon Carbide Transistor for Detecting Hydrocarbon Gases
- 25 Testing Thermophotovoltaic Cells With Black-Body Radiation
- 26 Swivel Total-Temperature Probe

Books and Reports

- 27 Thermal Conductance of Gold Contacts at Liquid-Helium Temperatures
- 28 Basis and Application of the CARES/LIFE Computer Program

BLANK PAGE

Integrated Oxygen-Recovery System

A reactor converts carbon dioxide and water to oxygen and methane without emitting hydrogen.

Marshall Space Flight Center,
Alabama

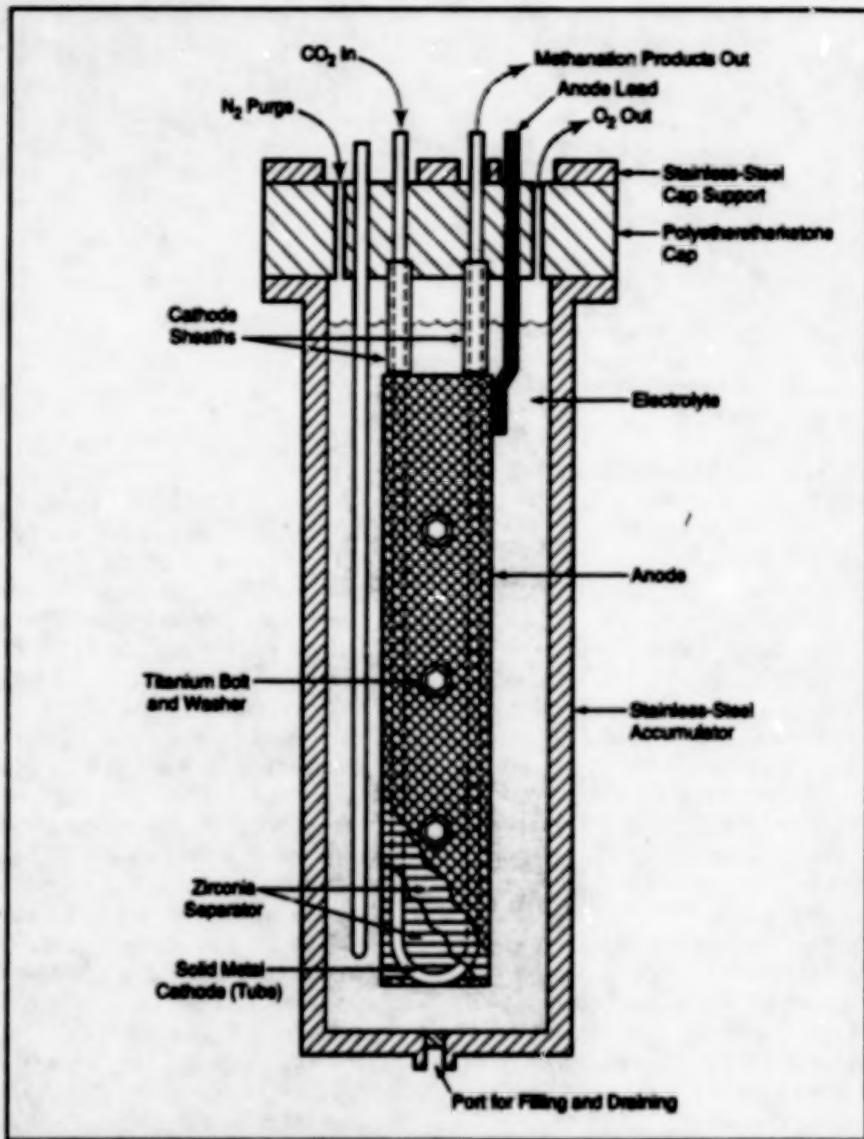


Figure 1. A Reactor Module in the integrated oxygen-recovery system features a simple, compact design. It contains both an electrolyzer and a Sabatier reactor.

A developmental system for revitalizing breathing air in an enclosed environment converts carbon dioxide and water to oxygen and methane. These chemical conversions take place in a reactor module which contains a solid-metal-cathode water electrolyzer integrated with a Sabatier reactor. According to design estimates, a fully developed version of the system sized to maintain breathable air for four persons would occupy 67 percent less volume, weigh 60 percent less, and consume 3.4 percent less power than would a system of the same capacity but in which the electrolyzer and Sabatier reactor would be distinct components. Although oxygen is the primary desired end product in the system as originally contemplated, the product

methane may also be of value in an industrial version of the system.

As in other systems, electrolysis of water in the reactor module breaks the water down into oxygen and hydrogen, the oxygen leaving the anode and the hydrogen leaving the cathode. The cathode is a tube that is part of an electrode assembly immersed in an aqueous electrolyte containing 55 percent KOH. The assembly also contains a porous anode that surrounds the cathode. A zirconia separator lies between the anode and the cathode (see Figure 1).

While still in monatomic form from the electrolysis, the hydrogen diffuses from the outside (electrolyte side) to the inside of the cathode tube, combining into hydrogen

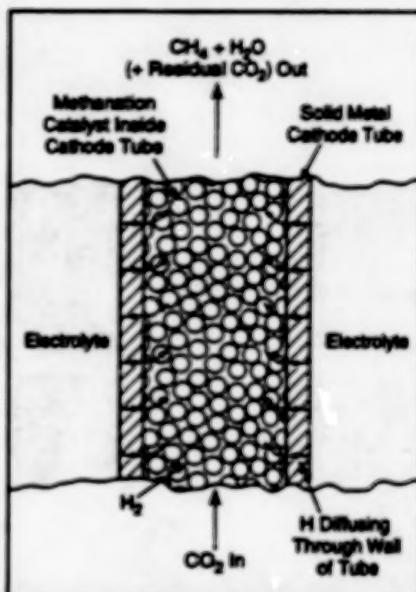


Figure 2. Hydrogen from electrolysis diffuses through the wall of the cathode into the interior space containing the methanation catalyst.

molecules at the interior surface. The hydrogen is used in the Sabatier reaction, in which carbon dioxide and hydrogen react in the presence of a catalyst to form water and methane. Heretofore, the Sabatier reactor would have been a separate unit.

The advantages of this system arise from the placement of the methanation catalyst inside the cathode tube instead of in a separate reactor unit. This enables transfer of the hydrogen directly and nearly instantaneously from one reaction to the other, so that little or no free hydrogen leaves the reactor. Another advantage of this system is that the water consumed in the electrolysis is reclaimed in the Sabatier reaction. Yet another advantage is that even the heat generated by the Sabatier reaction can be used to help maintain an elevated reactor temperature that facilitates electrolysis. This reduces the power consumed by the electrolyzer. Besides the production of oxygen, the methane produced may be of commercial value.

This work was done by F. H. Schubert, R. J. Davenport, and M. G. Lee of Life Systems, Inc., for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-26266.

Selective Oxidizer for Removal of Carbon Monoxide

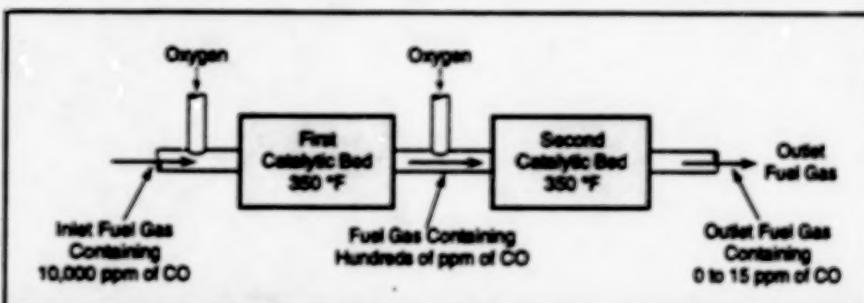
Multiple catalytic stages at progressively lower temperatures operate without becoming poisoned.

Lewis Research Center,
Cleveland, Ohio

A catalytic apparatus selectively oxidizes most of the carbon monoxide (without also oxidizing the hydrogen) in a stream of reformed fuel gas that is fed to a low-temperature fuel cell. It is necessary to oxidize most of the carbon monoxide because the CO content of a typical stream of reformed fuel gas is of the order of 1 percent, whereas a low-temperature fuel cell cannot tolerate CO in concentrations \geq 25 parts per million (ppm). Catalysts that can be used to oxidize CO selectively include platinum on alumina and a commercial catalyst known as "Selectoxo."

To provide for the oxidation of the CO, a small amount of oxygen or air is mixed into the stream of reformed fuel gas. The gas then flows through a series of catalytic beds. The beds are cooled under thermostatic control, such that each succeeding bed is maintained at a temperature lower than that of the preceding bed (see figure). The reasons for this staged operation at progressively lower temperatures are the following:

- A catalytic bed at a higher temperature [\sim 350 °F (177 °C)] typically reduces the level of CO levels from about 10,000 ppm (1 percent) to hundreds of ppm. Such a bed operates stably, with no increase in its outlet CO concentration with time. In other words, the catalyst in such a bed does not become poisoned by the CO.
- A catalytic bed at a lower temperature [\sim 170 °F (77 °C)] can initially reduce the concentration of CO from 10,000 ppm to 10 ppm. However, at this high inlet concentration of CO at the lower temperature, the bed is not stable in the



Catalytic Beds at Different Temperatures can oxidize most of the CO content in the inlet fuel gas stream for a long time without becoming poisoned. A single, lower-temperature bed would become poisoned, while a single, higher-temperature bed would not reduce the concentration of CO to an acceptably low level.

sense that its outlet concentration of CO increases with time. In other words, the catalyst becomes passivated; that is, poisoned by the CO.

The use of multiple catalytic beds at progressively lower temperatures eliminates the poisoning. The first and hottest bed receives the highest concentration of CO (10,000 ppm) and enough oxygen to oxidize the CO. Because this bed operates at high temperature, it is not poisoned. This hottest bed reduces the concentration of CO from 10,000 ppm to hundreds of ppm.

With its lower CO content, the stream of gas is then fed to a cooler bed, along with more oxygen. In the cooler bed, the CO content is reduced from hundreds of ppm to a level between 0 and 15 ppm. Because the gas that enters the second bed contains much less CO than does the gas that enters the first bed, the second bed is less susceptible to poisoning by CO and its life is extended accordingly.

This concept has been verified by an experiment on a prototype of the apparatus containing two catalytic beds. The first bed was operated at 350 °F (177 °C) with 1 percent of CO in the inlet gas. The second bed was operated at 170 °F (77 °C).

At the end of 330 h, the concentration of CO at the outlet from the first bed and the inlet for the second bed was 275 ppm, and the concentration of CO at the outlet of the second bed was 6 ppm. In contrast, the concentration of CO at the outlet of a single catalytic bed operating with an inlet CO concentration of 1 percent at a temperature of 170 °F (77 °C) increased from an initial level of 5 ppm to a level of 60 ppm after 135 h.

This work was done by John C. Trociale, Craig R. Schroll, and Roger R. Lissner of International Fuel Cells Corp. for Lewis Research Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Lewis Research Center [see page 1]. Refer to LEW-15630.

Silicon Carbide Transistor for Detecting Hydrocarbon Gases

A porous gate would allow passage of gas molecules.

A proposed silicon carbide variable-potential insulated-gate field-effect transistor would be specially designed for use in measuring concentrations of hydrocarbon gases. Devices like this one should prove useful in numerous automotive, industrial, aeronautical, and environmental monitoring applications.

The device (see figure) would include a sensing gate that would comprise an

open-metal-grid electrical contact (such as chromium) on a layer of porous undoped SiC that has been photo-electrochemically etched in a 6H-SiC wafer. The pores in the undoped SiC layer would be oriented perpendicular to the top of the device by the etching process. The sensing gate would be on an aluminum nitride insulating layer on top of a layer of n⁻-doped SiC mounted on a suitable dielectric support. Drain and source contacts would be attached to the n⁻ layer, and a variable potential contact would be attached to the metal grid.

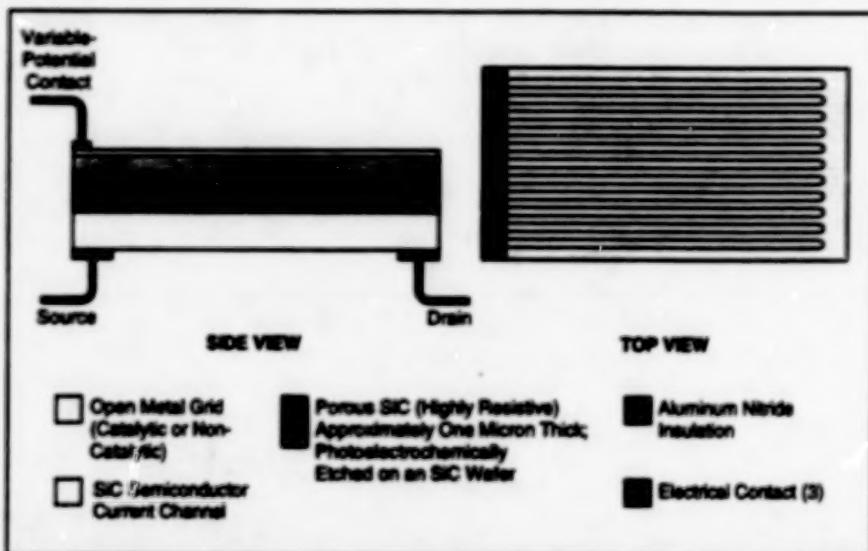
Hydrocarbon gas molecules would diffuse through the open metal grid into the porous SiC layer, where they would be adsorbed onto the surfaces of the pores. The hydrocarbon-gas sensor operates by dissociating or electrochemically oxidizing

NASA's Jet Propulsion Laboratory,
Pasadena, California

these hydrocarbons adsorbed to the inner surfaces of the SiC detection layer formed by the porous SiC. Dissociation or oxidation reactions are driven by a varying potential applied to the detection layer. Different hydrocarbon species undergo these reactions at different applied potentials so that the device is able to discriminate among various hydrocarbon species. Hydrogen ions produced by the dissociation reactions would travel to the surface of the aluminum nitride insulating layer, giving rise to an electric field that would, in turn, result in a depletion layer in the n⁺-doped SiC channel on the opposite side of the insulator. As in other field-effect transistors, the magnitude of this electric field and thus the extent of the depletion layer would affect the amount of current flowing between the source and drain.

The device can operate at temperatures between 100 °C and at least 600 °C, allowing hydrocarbon detection in hot exhaust gases. The dissociation reaction is detected as a change in the current flow through the transistor. The SiC detection layer can be augmented by replacing the metal grid with one composed of a catalytic metal, which provides a signal without an applied potential at higher operational temperatures. Comparisons between the catalytically produced signal and the varying-potential produced signal may further help identify the hydrocarbons present.

The magnitude of the variable potential on the porous SiC layer would be a measure of the strength of the molecular bond that could be broken by the applied electric field. Thus, to the extent to which the strengths of molecular bonds could be



Bond Breakage due to dissociation and electrochemical oxidation in the pores of the porous SiC layer would generate ions, which would migrate to the surface of the aluminum nitride. There, the electric field generated by the accumulated ions would affect the transistor current according to the basic principle of operation of a field-effect transistor.

used to distinguish among chemical species, the applied potential needed to obtain a nonzero reading could be taken as an indication of the detected molecular species. The magnitude of the current in the source-to-drain channel would indicate the concentration of the species.

The number of molecules adsorbed on the surfaces of the pores could eventually become excessive. In that case, the potential applied to the variable-potential contacts would be reversed and increased to drive off the adsorbed molecules.

A prototype sensor composed of the gate structure of the transistor has been fabricated and is undergoing tests. Preliminary results have demonstrated an ability

to discriminate between inert gases and hydrocarbon gases such as methane and propylene from about 100 °C to 400 °C.

This work was done by Virgil B. Shields, Margaret A. Ryan, and Roger M. Williams of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, NASA Resident Office-JPL [see page 1]. Refer to NPO-19428.

Testing Thermophotovoltaic Cells With Black-Body Radiation

Test conditions can be varied conveniently and are reproducible.

NASA's Jet Propulsion Laboratory, Pasadena, California

A method of testing thermophotovoltaic (TPV) cells involves maintaining the cells at various temperatures while exposing the cells to black-body radiation at various intensities from black-body sources at various other temperatures. This method provides for simulation and evaluation of performance under a variety of operating conditions and could thus become a standard testing method for government and industry.

In a TPV apparatus, heat is first converted into photons via a thermal emitter, and these photons in turn are converted to electricity via TPV cells, which are similar to ordinary photovoltaic cells except

that they are fabricated to have a bandgap that corresponds to the peak of the spectrum of radiation from the thermal emitter. In principle, the energy-conversion efficiency of such an apparatus can exceed that of a conventional photovoltaic cell if the apparatus comprises a thermal emitter of high emissivity and bandgap-matched TPV cells.

Heretofore, testing of TPV cells has involved measuring their current-vs.-voltage responses during irradiation from lamps or from selective emitters. Such tests have revealed little about performances of cells under various operating conditions because the test conditions

(emission temperature and/or emission spectrum, emission intensity, and cell temperature) have rarely been varied. Furthermore, such tests have not been reproducible because the emission temperatures and/or spectra have been incompletely characterized and highly specific to the thermal emitters used.

The present method overcomes these limitations, providing both variability and reproducibility. The black-body spectrum is characterized by Planck's equation and is easily reproducible with a standard black-body emitter. A black-body emitter can be set at various emission temperatures, thus altering the spectrum of the

emission as needed to test TPV cells. The intensity of black-body radiation incident on a TPV cell under test can be varied by simply placing the cell at various distances from the black-body emitter. Because the black-body spectrum is known, the results

of tests at various emission temperatures and intensities can be used to predict responses of TPV cells to radiation from non-black-body sources.

This work was done by James J. Lin and Dale R. Burger of Caltech for

NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].
NPO-19333

Swivel Total-Temperature Probe

A probe aligns itself with airflow to measure total temperature at high angle of attack.

Dryden Flight Research Center,
Edwards, California

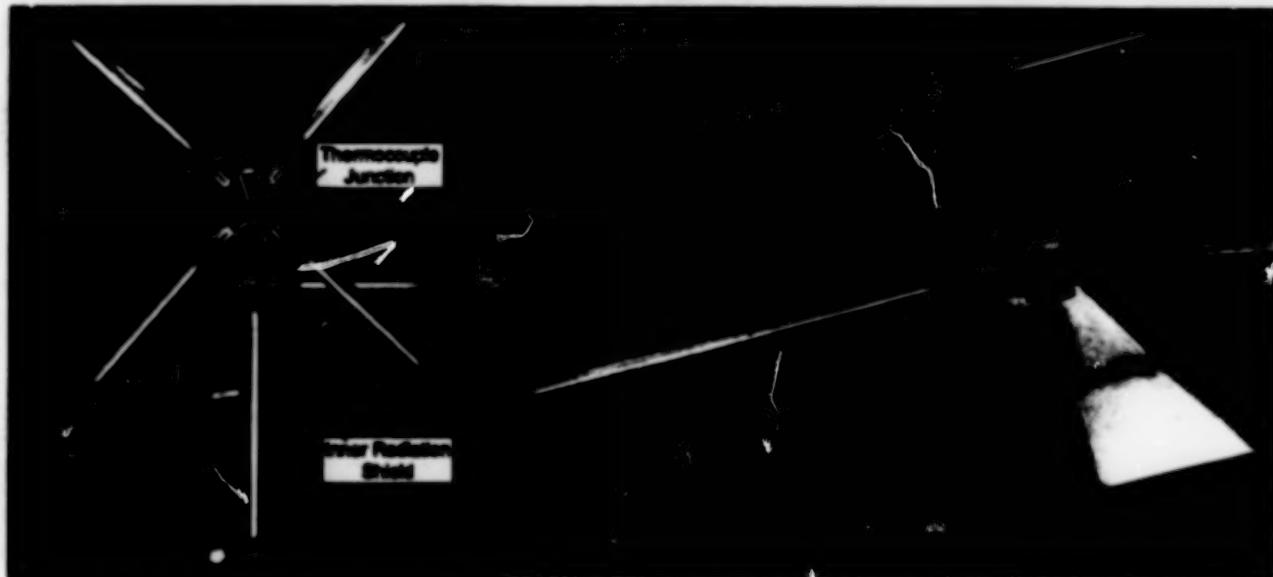


Figure 1. The Four Fins and Swivel Mount enable this probe to align itself aerodynamically with the local flow of air. Thus, the stagnated-flow condition needed for accurate measurement of total temperature can be achieved.

Figure 1 illustrates a swivel total-temperature probe. This is an improved probe designed to be mounted at the wingtip of an airplane and to be used for measuring the total temperature of the air when the airplane flies at high angles of attack.

Flight at high angles of attack has recently become more popular in the fighter community for its potential in tactical maneuvering. Current total-temperature probes on fighter aircraft were not designed for flight at high angles of attack. A total-temperature measurement is combined with a mach number to compute the ambient temperature, which is necessary for computing the true airspeed and Reynolds number. It is questionable whether total-temperature measurements at high angles of attack will be necessary on production fighters because control laws and flight-operation decisions are typically based on mach number, calibrated airspeed, or equivalent airspeed; none of these

quantities depends on measurement of total temperature. However, true airspeed and Reynolds number are of interest to the flight-research community. Consequently, there was a need to develop a total-temperature probe for flight research at high angles of attack.

A standard total-temperature probe is potentially subject to large errors at high angles of attack because (1) its orientation is fixed and thus cannot be adjusted to stagnate the flow (as is necessary for measuring total temperature) at high angles of attack and (2) deicing heat at the low airspeeds associated with high angles of attack can distort probe readings. When using a fixed total-temperature probe, errors of the first-mentioned type can be mitigated by mounting the probe on the underside of the fuselage to take advantage of a flow-straightening effect of the fuselage at high angle of attack, while errors of the second-mentioned type can be mitigated only by turning off deicing heaters.

The swivel probe, which does not contain deicing heaters, was designed to align itself aerodynamically with the local airflow. The swivel total-temperature probe includes a type-K (Chromel/Alumel) thermocouple within an inner radiation shield embedded in a tube similar to a pitot-tube on a wingtip-mounted air-data boom.

The performance of the swivel total-temperature probe has been evaluated in flight tests at the NASA Dryden Flight Research Center, using the F-18 High Alpha Research Vehicle (HARV). Excellent results have been achieved in dynamic flight at angles of attack as large as 70°. At the time of submission of information for this article, seventy flights had been completed, with no structural or instrumentation failures.

The data plotted in Figure 2 include results for a fairly dynamic high-angle-of-attack pushover/pullup maneuver. Ambient-temperature profiles from a fuselage-mounted fixed probe, the swiv-

el probe, and a standard-day calculation are shown. The swivel-probe data have been adjusted for an instrumentation bias, the cause of which has not yet been ascertained. The data obtained from the swivel total-temperature probe follow the trends of the standard-day temperature profile much more closely than do the data from the fixed probe. The errors in the fixed-probe data include the high-angle-of-attack and deicing-heater errors mentioned above. Inasmuch as the flow-straightening effect of the F-18 fuselage mitigates much of the high-angle-of-attack error, a major part of the fixed-probe error is attributed to deicing heat. Also shown in Figure 2 is a plot of the difference between the true airspeeds computed from temperature readings of the fixed and swivel probes during the maneuver in question. The magnitude of the difference was as much as 25 ft/s (7.6 m/s), which is approximately 8 percent of the true airspeed.

These results are preliminary. Work is continuing in an attempt to (1) understand the cause of the bias in the swivel-probe temperature measurements and (2) document the accuracy of the swivel-probe measurements at numerous static and dynamic flight conditions that involve high angles of attack. The fixed total-temperature probe works remarkably well at high angles of attack with the deicing heater turned off; if such operation is allowable, flight-research data of acceptably high quality can be obtained by use of a fuselage-mounted fixed total-temperature probe.

This work was done by Tim Moes, Doug Nichols, Dirk Vanderloop, and Rich Rood of Dryden Flight Research Center. For further information, make contact with Tim Moes by electronic mail at moes@wibur.dfrc.nasa.gov. DRC-95-21

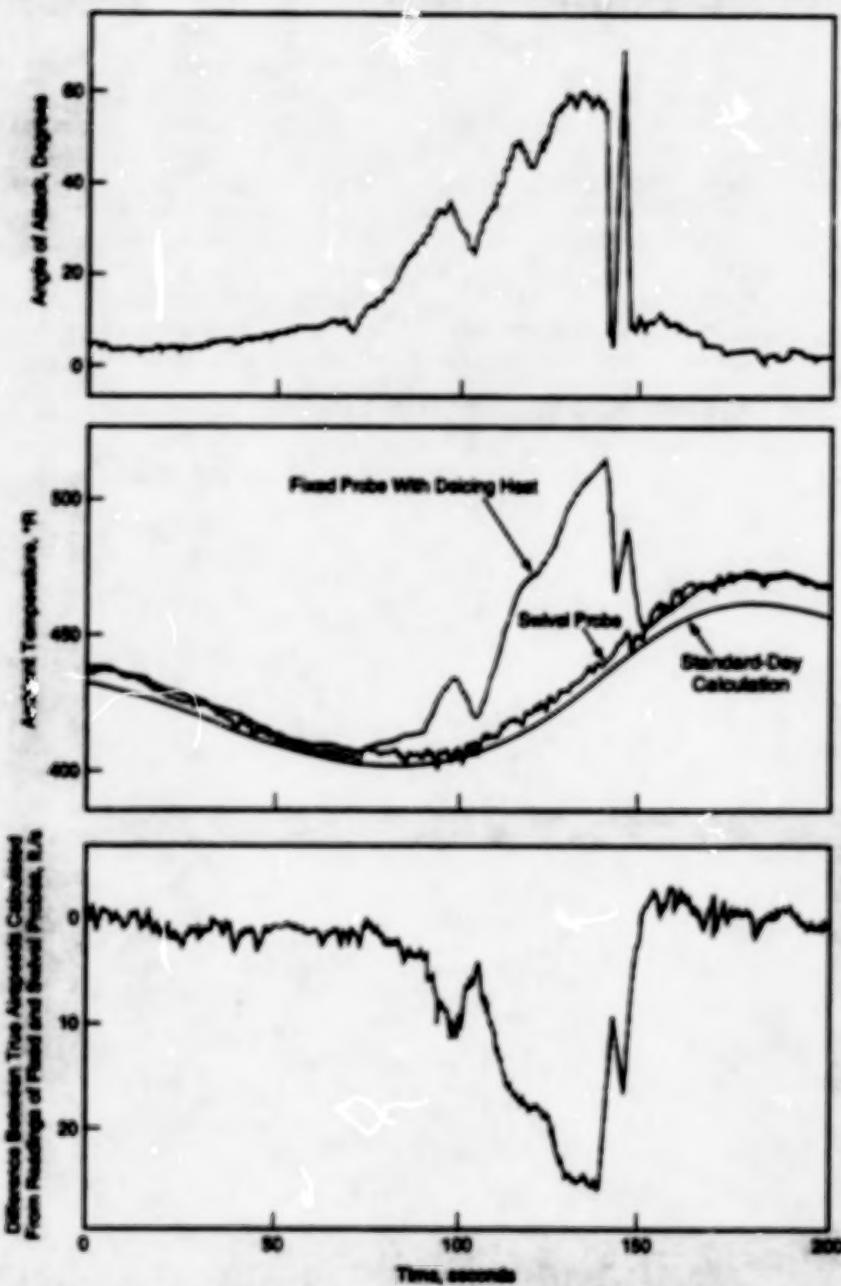


Figure 2. These Data Were Obtained in a flight test that involved highly dynamic maneuvering at high angles of attack.

Books and Reports

Thermal Conductance of Gold Contacts at Liquid-Helium Temperatures

A report describes a study of the thermal conductance of gold-plated copper, aluminum, brass, and stainless-steel contacts at temperatures from 1.6 to 4.2 K under applied forces of 22 to 670 N. It

was found that the thermal conductance improved greatly as the result of gold-plating, except for the stainless-steel samples. All samples were prepared by first degreasing, then ion-milling, followed by a deposition of a 100-nm-layer of chromium and 2 μm of gold.

This work was done by Peter Kötter and Louis J. Salerno of Ames Research Center and Alan L. Spivak of Trans-Bay

Electronics. To obtain a copy of the report, "Thermal Conductance of Gold-Plated Metallic Contacts At Liquid Helium Temperatures," see TSP's [page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Ames Research Center [see page 1]. Refer to ARC-13178.

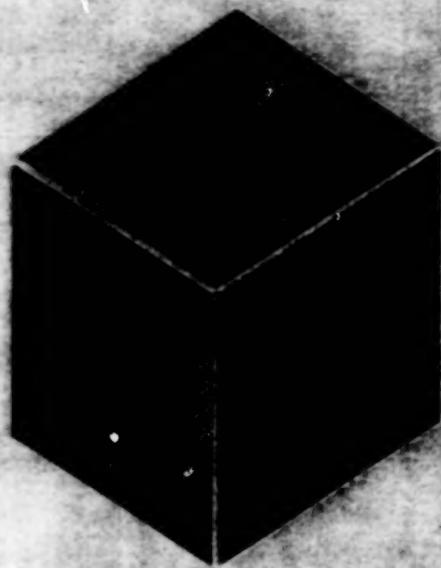
Basis and Application of the CARES/LIFE Computer Program

A report discusses the physical and mathematical basis of the Ceramics Analysis and Reliability Evaluation of Structures LIFE prediction (CARES/LIFE) computer program, which was described in "Program for Evaluation of Reliability of Ceramic Parts" (LEW-16018), page 43. As explained in more detail in that article and in the report, this software calculates the time-dependent reliability (probability of failure as a function of time in service) of a monolithic ceramic component subjected to thermomechanical and/or proof-test

loading. The report explains the basic equations implemented in CARES/LIFE to mathematically model the various phenomena involved in the failures of ceramic components: CARES/LIFE accounts for subcritical crack growth by use of a power law, the Paris law, or the Walker equation. The two-parameter Weibull cumulative distribution function is used to characterize the variation in the strengths of components. The effects of multiaxial stresses are modeled by use of either the principle of independent action, the Weibull normal-stress-averaging method, or the Batdorf theory. Inert strength and fatigue parameters are estimated from rupture-strength data from naturally flawed specimens

loaded in static, dynamic, or cyclic fatigue. Following these explanations, the report describes a test case in which CARES/LIFE was used to predict time-dependent reliabilities of alumina bar and disk flexure specimens that exhibited subcritical crack growth when exposed to water.

This work was done by Noel N. Nemeth, Lesley A. Janosik, and John P. Gyekenyesi of Lewis Research Center and Lynn M. Powers of Cleveland State University. To obtain a copy of the report, "Durability Evaluation of Ceramic Components Using CARES/LIFE," see TSP's [page 1]. LEW-16207



Materials

Hardware, Techniques, and Processes

- 31 Skutterudite Compounds for Power Semiconductor Devices
- 32 Using Aerospace Technology To Design Orthopedic Implants
- 33 Multidisciplinary Design of Hot Composite Structures
- 35 Computer-Aided Process Model for Carbon/Phenolic Materials
- 35 Epitaxial Silicon Doped With Antimony
- 36 CVD of Thin Films From Single Organometallic Precursors
- 36 Modified Process for Formation of Silicon Carbide Matrix Composites

BLANK PAGE

Skutterudite Compounds for Power Semiconductor Devices

These materials exhibit high hole mobilities, high doping levels, and high electronic figures of merit.

NASA's Jet Propulsion Laboratory,
Pasadena, California

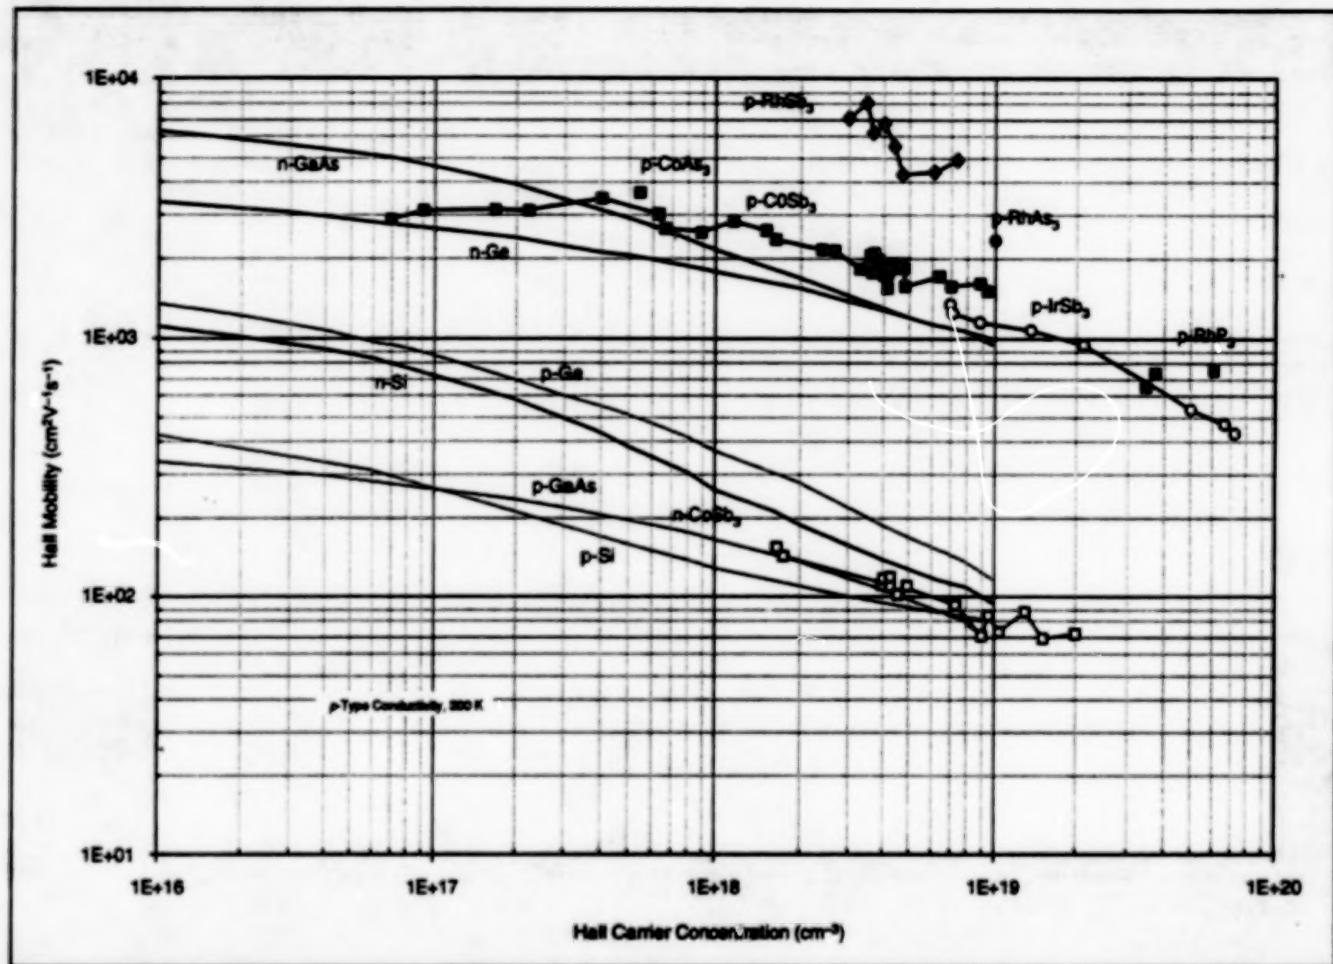


Figure 1. In Hall Mobility Versus Carrier Concentration at 300 K for p-type skutterudite compounds, data are compared to n-type and p-type state-of-the-art electronic materials. (n-type CoSb₃ data are also included.)

New semiconducting materials with p-type carrier mobility values much higher than those of state-of-the-art semiconductors have been discovered. The nine compounds, antimonides CoSb₃, RhSb₃, IrSb₃, arsenides CoAs₃, RhAs₃, IrAs₃, and phosphides CoP₃, RhP₃ and IrP₃, exhibit the same skutterudite crystallographic structure and form solid solutions of general composition $\text{Co}_{1-x}\text{Rh}_x\text{Ir}_y\text{P}_{1-y}\text{As}_z\text{Sb}_2$. Some of these compositions have shown great potential for application to thermoelectric devices.

There are also a number of more complex skutterudite compositions including solid solutions and ternary compounds. This family of semiconducting materials offers great potential for utility in (a) high-frequency power-semiconductor devices like power field-effect transistors, (b) these

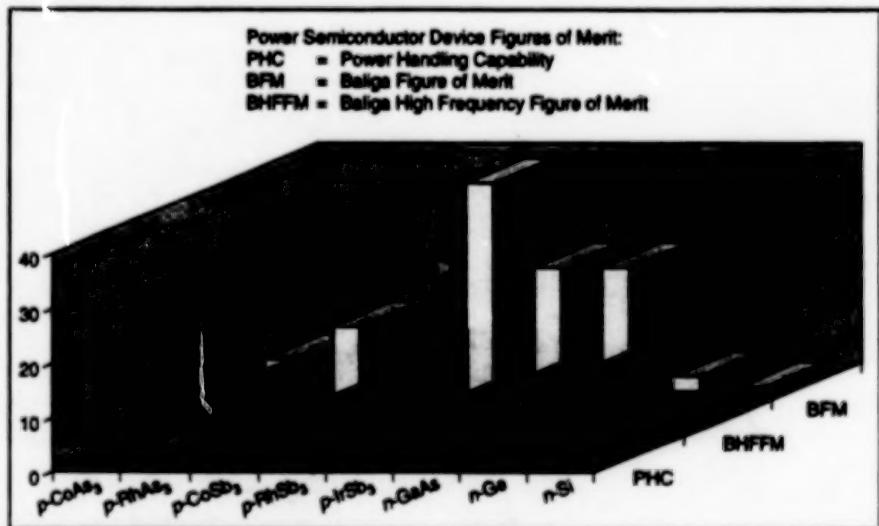


Figure 2. Figures of Merit are shown for p-type skutterudite and state-of-the-art electronic materials for power semiconductor devices. (All data for a carrier concentration of $1 \times 10^{19} \text{ cm}^{-3}$ are normalized to n-type Si.)

and other electronic devices at temperatures above 150 °C where Si cannot be used, and (c) other electronic devices (e.g., tunnel diodes and lasers) made of heavily doped materials. These skutterudite compounds are relatively easy to prepare, and some of them are reasonably inexpensive. Provided lower carrier concentrations can be achieved for these new materials, they might even challenge current state-of-the-art lightly-doped electronic materials used in a variety of junction devices.

The antimonides CoSb₃, RhSb₃, IrSb₃, and arsenides CoAs₃ and RhAs₃ all exhibit semiconducting behavior, with band-gap values ranging from 0.63 to 1.18 eV. The phosphide compounds, currently under investigation, are expected to have even wider band gaps. Due to the exceptionally high hole mobilities, *p*-type skutterudites exhibit high electrical-conductivity values ranging from 2 to $5 \times 10^5 \Omega^{-1} \text{ m}^{-1}$ for a hole concentration of $1 \times 10^{25} \text{ m}^{-3}$.

The room temperature mobility values of *p*-type skutterudites are 1 to 100 times higher than that for *p*-type Si and GaAs at similar carrier concentrations, as seen in the first figure. RhSb₃ exhibits the greatest hole mobility, $8000 \times 10^{-4} \text{ m}^2 \text{ V}^{-1} \text{ s}^{-1}$ for a carrier concentration of $2.5 \times 10^{24} \text{ m}^{-3}$.

which is about 70 times higher than *p*-type GaAs and still 5 times higher than *n*-type GaAs. Skutterudites with *n*-type conductivity can be obtained by doping with selected elements. The electron mobilities of *n*-type CoSb₃ are comparable to the values obtained for *p*-type Si and GaAs. Due to the preparation techniques used thus far, skutterudite samples with carrier concentration $< 7.0 \times 10^{22} \text{ cm}^{-3}$ have not yet been obtained, but it is expected that the carrier mobility will increase with decreasing carrier concentrations. The temperature dependencies of the carrier mobilities of the skutterudites also compare favorably with those of state-of-the-art semiconductors. Even at 550 °C, the hole mobility of IrSb₃ has been observed to be about $600 \times 10^{-4} \text{ m}^2 \text{ V}^{-1} \text{ s}^{-1}$ for a carrier concentration of $6.5 \times 10^{24} \text{ m}^{-3}$ —about 4 times the value obtained for *n*-type Si.

The second figure shows the power semiconductor figures of merit for *p*-type skutterudites and *n*-type state-of-the-art semiconductors Ge, Si, and GaAs calculated for a carrier concentration of $1 \times 10^{25} \text{ m}^{-3}$. This graph demonstrates that the high mobilities at high doping levels of the *p*-type skutterudites more than offset their low thermal-conductivity values (one

fourth of GaAs), even when comparing to the high electron mobility *n*-type state-of-the-art electronic materials. This means that the temperature rise of junctions above ambient temperature will be negligible, minimizing heat dissipation problems. Also, the mobilities of the skutterudites decrease slowly with temperature, and this raises the possibility of operating electronic devices at temperatures much higher than possible when using devices made of Si.

This work was done by Jean-Pierre Fleurial, Thierry Caillat, Alexander Borshchevsky, and Jan Vandersande of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].

In accordance with Public Law 96-517, the contractor has elected to retain title to this invention. Inquiries concerning rights for its commercial use should be addressed to

William T. Calleghan, Manager
Technology Commercialization
JPL-301-350
4800 Oak Grove Drive
Pasadena, CA 91109

Refer to NPO-19378, volume and number of this NASA Tech Briefs issue, and the page number.

Using Aerospace Technology To Design Orthopedic Implants

Methods for optimizing designs of turboprop blades have been applied to prosthetic knee joints.

Technology originally developed to optimize the designs of composite-material aerospace structural components has found a new use in efforts to develop a method for optimizing the designs of orthopedic implants. Thus far, the development effort has focused on designing knee implants, but the long-term goal of this effort is to develop a method for optimizing the designs of orthopedic implants in general.

The prototype method provides for tailoring the shapes of the tibial components of a total-knee-replacement implant for optimal interaction within the environment of the tibia. The shapes of the components of the implant are optimized such that the stresses in the bone are distributed more favorably than they otherwise would be, thereby minimizing the degradation of bone, improving the mechanical integrity of the bone/implant system, and preventing failures of the components of the implant.

The computer code used in the prototype method is an adapted version of the Structural Tailoring of Advanced Turboprops (STAT) computer code. Despite the

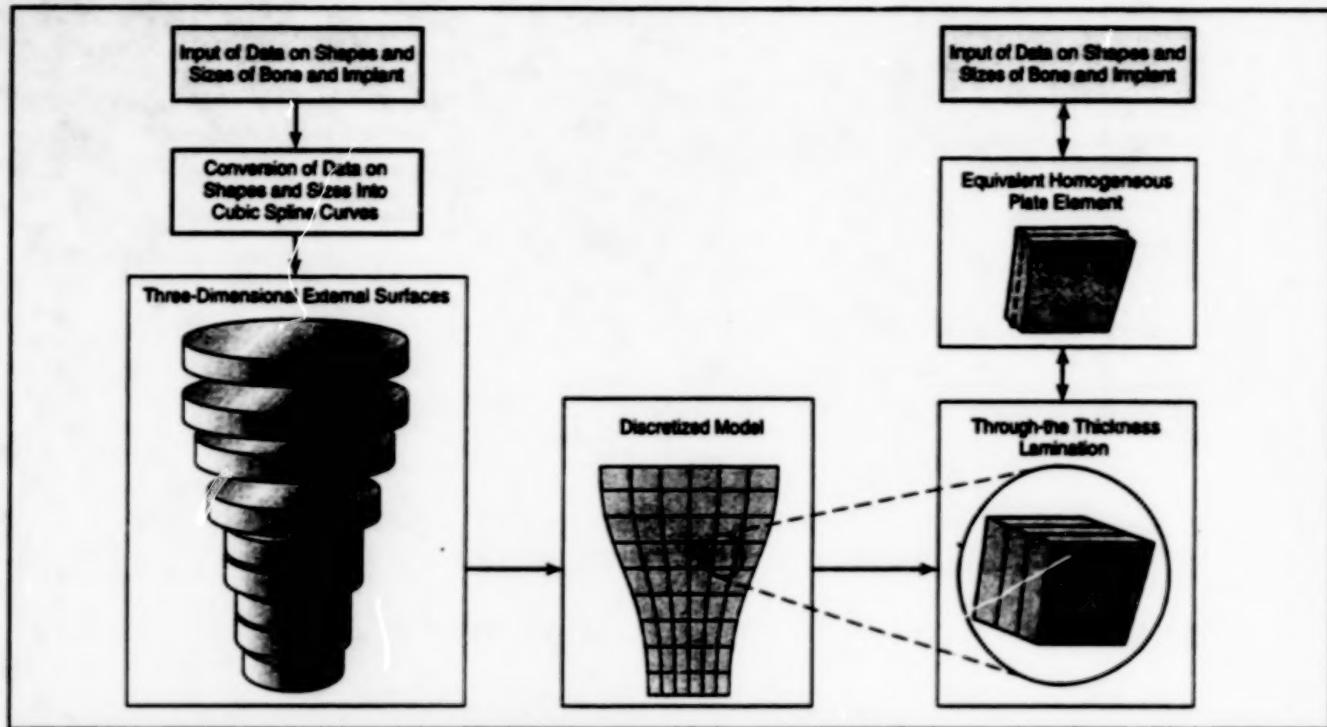
obvious differences between laminated composite turboprop blades on the one hand and bone/implant systems on the other hand, the essential features of STAT are applicable to both; these features include composite-mechanics, finite-element-analysis, and structural-optimization capabilities. An essential part of the optimization procedure is mathematical modeling of the quasi-static mechanical response of the laminated tissue structure of the upper end of the tibia. This modeling involves utilization of the finite-element-modeling capability of STAT and, in particular, laminated-plate models (see figure) as a compromise between overly simplistic two-dimensional models and fully complex three-dimensional models.

The analysis of stresses and strains takes place in three stages. In the first stage, stiffness matrices that represent the characteristics of the discrete bone layers are synthesized. Integration through the sagittal plane (thickness) provides the equivalent local through-the-thickness extensional and flexural stiffness matrices

Lewis Research Center,
Cleveland, Ohio

of the laminate, relating the generalized (average) laminate stresses to the generalized (average) laminate strains. In the second stage, the global static response is obtained via finite-element analysis with discretized nodal displacements in the frontal plane. The third and final stage involves the back-calculation of stresses within the bone and the implant components from the calculated global structural response. The generalized strains in the frontal plane are calculated, then the strain in each discretized layer is calculated. The stresses in the discretized material layers are then calculated from the layer strains.

The design parameters subject to optimization include those that specify the shape of a post and the thickness of a metal backing tray that are parts of the implant. The basic criterion for optimization of these parameters is assumed to be an acceptable state of stress in the bone/implant composite structure. For the tibial component, the most likely mode of failure is aseptic loosening, which involves the resorption of the bone adjacent to the



Laminated-Plate Models are used in the initial finite-element analysis of the bone/implant system.

implant. This is attributed in many cases to the stress-induced adaptation of the bone after the implantation surgery in response to the altered stress field induced exclusively by the presence of the implant. Thus, the optimal-design goal becomes the minimization of undesirable atrophy of the adjacent bone, or, in a positive sense, the creation of a favorable stress field in the bone.

Accordingly, the optimization-and-design procedure is an iterative one in which the shape of the post and the

thickness of the metal backing tray are made to evolve in small increments, after each of which a finite-element analysis is performed to determine the new stress field. That information obtained at each such step of the design can then be used to direct the next design move in an effort to reach the optimum design by minimizing an objective function. The objective function chosen initially was the maximum specific distortion energy near the bone/implant interface.

This work was done by D. A. Saravacos, P. J. Mraz, and D. T. Davy of Case Western Reserve University for Lewis Research Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Lewis Research Center [see page 1]. Refer to LEW-16093.

Multidisciplinary Design of Hot Composite Structures

A computer code incorporates modules for simulation and analysis by traditionally separate disciplines.

A unified computer code has been developed to implement a multidisciplinary approach to the design and analysis of composite-material structures that must withstand high temperatures. The code is modular: it includes an executive module that communicates with and coordinates other modules that perform calculations pertaining to traditionally separate disciplines like those of acoustics, structural vibrations, structural loads, and thermal effects (see Figure 1).

An essential feature of the multidisciplinary approach is finite-element numerical simulation of the relevant physical phenomena according to the applicable disciplines. The same finite-element mesh is used in the thermal, vibrational, and struc-

tural analyses; this minimizes data-preparation time and eliminates the errors incurred in transforming temperatures from one finite-element mesh to another.

The integrated multidisciplinary simulation procedure is founded, in part, on recognition of the essential complexity and nonlinearity of the coupled heat-transfer, geometrical, material-property, structural, acoustic and nonacoustic vibrational, and environmental effects. For example, the thermal, mechanical, and acoustic properties of a composite material are nonlinear functions of temperature and other environmental parameters. Accordingly, the unified program includes a module that contains the Integrated Composite Analyzer (ICA) computer program, which

Lewis Research Center,
Cleveland, Ohio

implements micromechanics- and laminate-oriented mathematical models of composite materials to compute both the macroscale and the local microscale hygro/thermo/mechanical responses of composite materials, including effects of deterioration of properties with time.

Figure 2 illustrates the multidisciplinary design-tailoring and analysis procedures. The first step in the design-tailoring procedure is the specification of the tailoring input, which consists of an objective function, design variables, constraints, and convergence tolerances. The next step is to establish the current design via an initial multidisciplinary analysis based on values of the current design variables. The mathematical representation of the

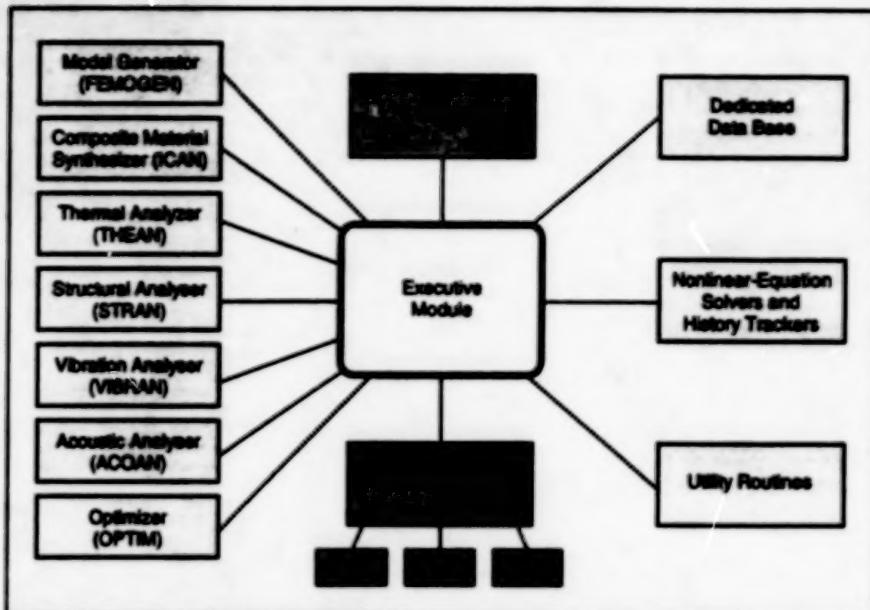


Figure 1. A Unified Computer Code includes modules for simulation and analysis by traditionally separate disciplines.

discipline-specific loads (e.g., thermal and mechanical loads) activates the cor-

responding multidisciplinary analysis modules.

The rest of the tailoring procedure is iterative. The gradients of the objective function and constraints are computed with respect to the design variables. A search vector is defined. A new set of design variables is chosen along the direction of the search vector. Again using the multidisciplinary analysis module, a proposed design based on the new design variables is then formulated and the physical behavior of the structure with the proposed design is simulated numerically. Then the objective function, the design variables, and the constraints are checked, on the basis of specified tolerances, to determine whether the proposed design is optimum. The procedure is repeated until an optimum design that satisfies all the constraints is achieved.

This work was done by Christos C. Chemis of Lewis Research Center and Surendra N. Singhal of NYMA. Further information is contained in a TSP [see page 1]. LEW-15977

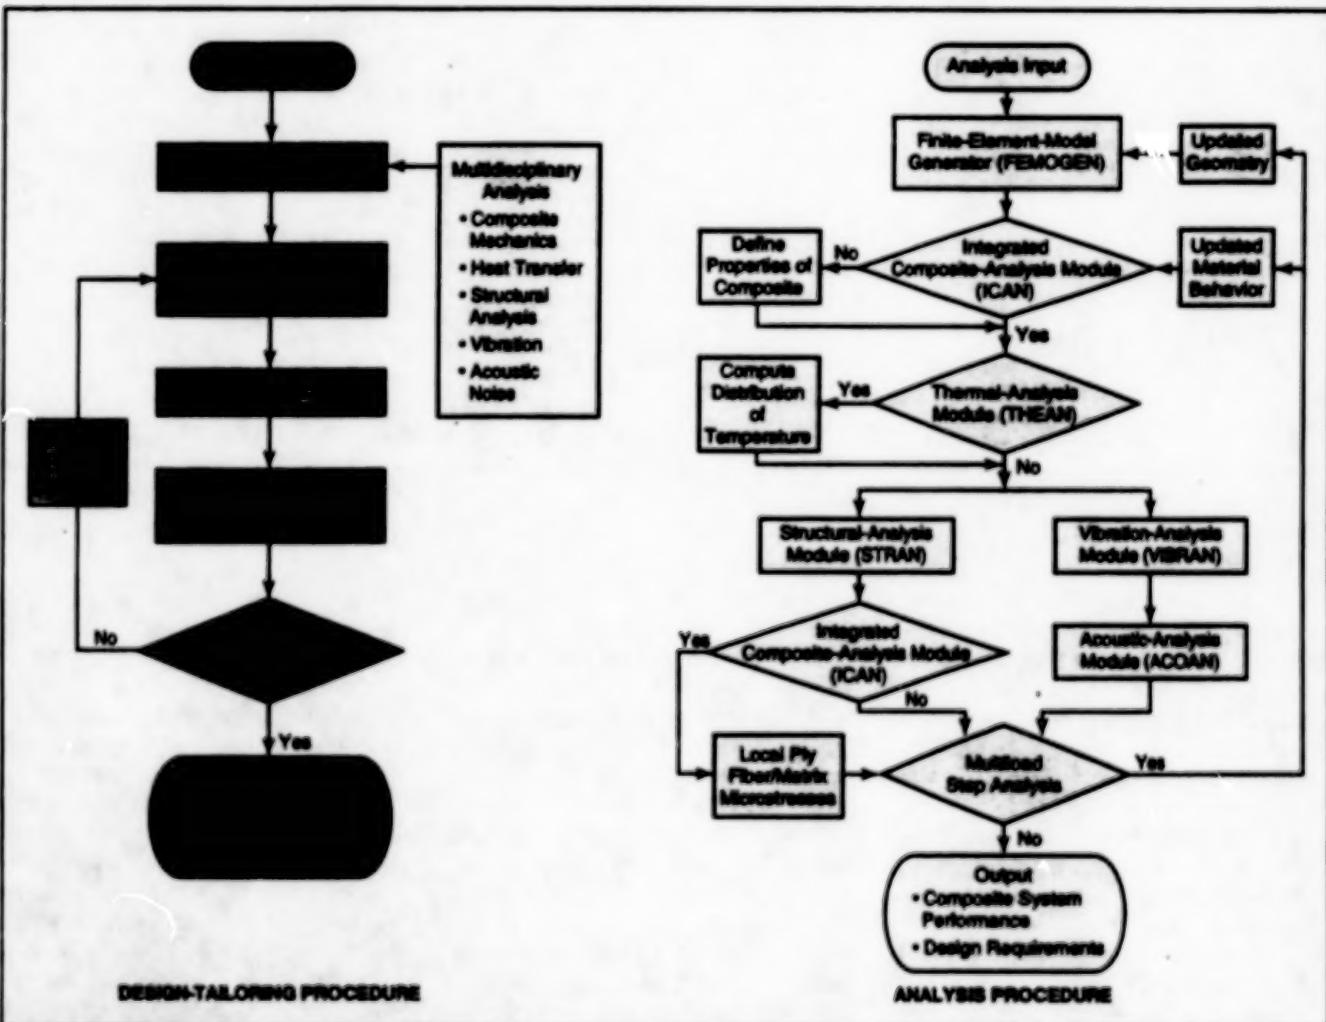


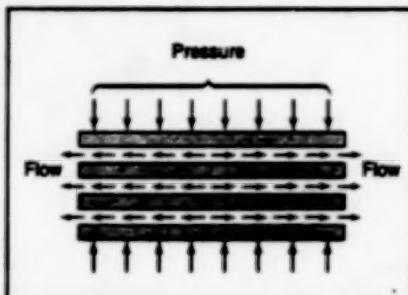
Figure 2. Multidisciplinary Design and Analysis Procedures that are implemented in the computer code reflect the nonlinearity and complexity of the relevant physical phenomena and of the design process.

Computer-Aided Process Model for Carbon/Phenolic Materials

The model can be used to optimize curing processes and predict the properties of the materials.

A computer program implements a thermochemical model of the processing of carbon-fiber/phenolic-matrix composite materials into molded parts of various sizes and shapes. The program and model are subjects of a continuing development effort. Although this effort is directed toward improving the fabrication of rocket-engine-nozzle parts made of carbon/phenolic materials, the model can also be used to optimize the fabrication of other structural components, and its material-property parameters can be changed to apply to other materials. The model helps to reduce costs by reducing the amount of laboratory trial and error needed to optimize curing processes and to predict the properties of the cured parts.

The program and model are based on the SINDA heat-transfer program and model. The program incorporates sophisticated computer-graphical capabilities; it can provide a three-dimensional view of a part being analyzed, the position of the part in the nozzle or other overall structure, and



Some of the Matrix Resin Flows out from between plies when a fiber/matrix composite laminate is squeezed during the curing process.

cross-sectional views of changes in temperature, degree of cure, moisture, and viscosity during the curing process. The model includes submodels of the transfer of heat, thermochemical reactions, generation and diffusion of moisture, viscosity, and the degree of cure. These submodels have been verified in laboratory tests.

A submodel of the flow of resin under pressure (see figure) is planned. This submodel will be based on data from flow

Marshall Space Flight Center,
Alabama

tests in a laboratory. The model has been verified by correlation with data on a commercial carbon-cloth/phenolic composite. With suitable adjustments of kinetics parameters, the predictions of the model also seem to be well correlated with data on some other composite materials.

Several enhancements of the program and model are planned. The first would be modification of the program to make it easier to use; anyone who has minimal knowledge of SINDA but at least some knowledge of computers would be able to use the model. The second enhancement would be completion and verification of the resin-flow submodel. The third enhancement would be the addition of a submodel and code for the calculation of residual stresses and of compaction during cure.

This work was done by Mischell A. Letson and Robert C. Bunker of Thiokol Corp. for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].
MFS-28969

Epitaxial Silicon Doped With Antimony

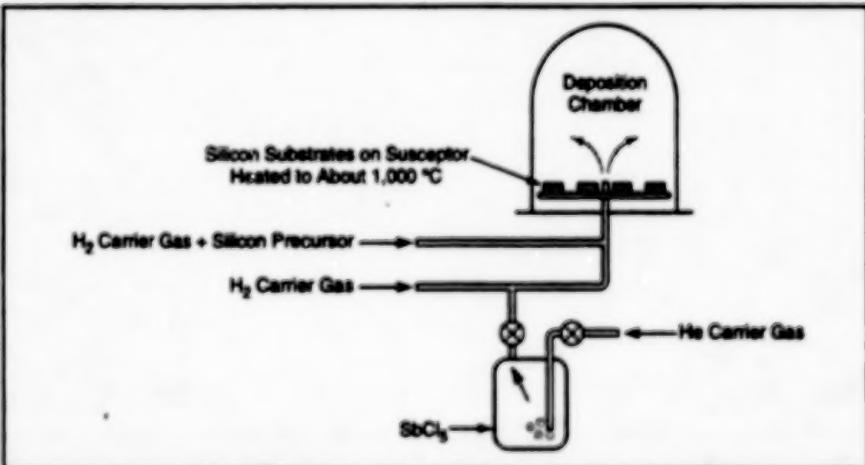
High purity can be achieved in layers of arbitrary thickness.

High-purity epitaxial silicon doped with antimony can be made by chemical vapor deposition, using antimony pentachloride ($SbCl_5$) as the source of dopant and SiH_4 , $SiCl_2H_2$, or another conventional source of silicon (see figure). Epitaxial silicon doped with antimony is needed to fabricate impurity-band-conduction photodetectors that operate at wavelengths from 2.5 to 40 μm ; in contrast, impurity-band-conduction photodetectors previously made from silicon doped with arsenic operate in the narrower wavelength range of 2.5 to only 28 μm .

Previously, in chemical vapor deposition in which organometallic compounds were used as sources of dopants, it was found that the carbon in those compounds gave rise to carbon contamination in the deposited materials. The absence of carbon in the compounds used in the present method eliminates that particular source of contamination.

Epitaxial silicon doped with antimony can be deposited to arbitrary thickness by the present method. Demonstration deposits have been made in thicknesses

NASA's Jet Propulsion Laboratory,
Pasadena, California



$SbCl_5$ is the Source of Antimony Dopant in silicon deposited from chemical vapors. $SbCl_5$ is a liquid at room temperature and can be used in a typical chemical-vapor-deposition reactor like those in common use.

up to 25 μm , with residual impurity concentrations as low as 10^{12} atoms/cm 3 and with antimony-dopant densities from less than 10^{17} to more than 4×10^{18} atoms/cm 3 .

This work was done by James E. Huffman and Bradley L. Haileck of

Rockwell International Corp., Science Center for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].
NPO-19425

CVD of Thin Films From Single Organometallic Precursors

These films are superior to films deposited from multiple precursors.

Lewis Research Center,
Cleveland, Ohio

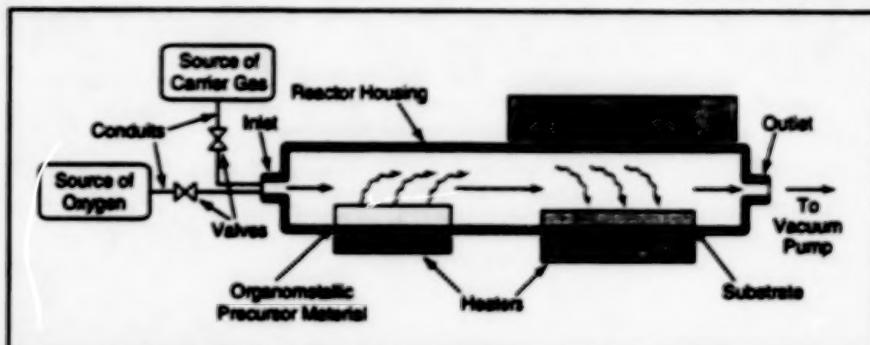
A method of forming thin inorganic films involves chemical vapor deposition (CVD) from single organometallic precursors. The method is especially suitable for depositing high-quality passivating or buffer layers of GaS on GaAs semiconductor substrates. The method is also applicable to formation of high-quality films for purposes other than buffering or passivation, and to different materials in which another element from the same group in the periodic table of elements is substituted for all or a portion of each element in the GaS/GaAs system.

Previous methods of forming passivating or buffer layers on GaAs involved multiple chemical precursors. These methods resulted, variously, in layers that contained impurities, were of uneven composition, were unstable, were toxic, and/or had to be deposited at temperatures so high that substrates became degraded.

In contrast, the present method involves no toxic constituents, minimizes impurities, and yields films that have substantially uniform crystal structure and composition.

The apparatus used in this method (see figure) includes a flow-through reactor housing connected via valves to (1) a vacuum pump, (2) a source of carrier gas (typically, argon, hydrogen, or nitrogen), and (3) an optional source of oxygen or a gas that can be decomposed to release oxygen (e.g., N_2O). The precursor material is placed in a holder in the chamber upstream from the substrate.

A heater plate outside the chamber near the precursor heats the precursor, through the wall of the reactor housing, to a temperature high enough to vaporize a sufficient quantity of the precursor; typically, this temperature is between 100 and 250 °C. A flow of carrier gas entrains the



The Organometallic Precursor Material is vaporized, carried to the substrate, and pyrolyzed to form a deposit on the substrate.

precursor and carries it toward the substrate. Another heater plate outside the reactor housing heats the substrate to a temperature at which the precursor is pyrolyzed; typically, this temperature is between 250 and 550 °C.

Upon arrival in the vicinity of the substrate, the precursor molecules are pyrolyzed by heat radiated from the substrate. The nonvolatile product (e.g., GaS) of the decomposition of the precursor molecules is deposited on the substrate, forming the film. The heaters and the flow of gas are turned off to stop the deposition when the deposit attains the desired thickness (typically between 2,000 and 5,000 Å).

The precursor can be any of a number of organometallic compounds that contain butyl groups and the elements to be deposited. The composition and molecular structure of the deposited film depends on the precursor. For example, a precursor of $[(Bu)_2Ga(SBu)]_2$ (where "Bu" denotes tertiary butyl) yields a film of amorphous GaAs.

An important class of examples is that of precursors that have cubane molecular structures or that acquire cubane molec-

ular structures during vaporization or pyrolysis. These precursors yield films that have cubic or cubiclike crystalline structures, which are often desired because they can be lattice-matched with substrates of similar crystalline structure. For example, the cubane-structured precursor $[(Bu)GaS]_4$ forms a cubic-phase GaS passivating film that is lattice-matched to a GaAs substrate.

The optional source of oxygen is for use in forming an insulating oxide layer on the outer surface of the passivating film. When the film has been deposited to the desired thickness, the valve from the source of oxygen is opened; exposure to the flowing oxygen or oxygen-containing gas results in formation of the oxide surface layer.

This work was done by Aloysius F. Hepp of Lewis Research Center; Andrew R. Baron of Harvard University; Michael B. Power and Andrew N. MacInnes of Galla, Inc.; and Philip P. Jenkins of Sverdrup Technology, Inc. Further information is contained in a TSP [see page 1].

LEW-15636

Modified Process for Formation of Silicon Carbide Matrix Composites

The modification reduces damage to SiC fibers during formation of an SiC matrix.

Lewis Research Center,
Cleveland, Ohio

A modified version of a process for making SiC-fiber/SiC-matrix composite material reduces damage to the SiC (SCS-6) fibers and to the carbon-rich coatings on the fibers. In both the unmodified and modified versions of the process, the matrix material that surrounds the fibers is reaction-formed silicon carbide (RFSC); it

is made by a succession of infiltration and pyrolysis steps, as explained more fully below. Ideally, the carbon-rich surface layers on the fibers should be retained as low-shear-strength interfacial layers that make the composite less brittle. In the unmodified process, these carbon-rich layers are destroyed and the underlying

SiC fibers are damaged, also as explained more fully below.

In both the unmodified and modified versions of the process, the first step is to infiltrate a stack of SiC fiber mats with a liquid polymer-forming mixture that comprises a high-char-yield monomeric resin, a pore-forming agent, a curing agent,

and carbon powder filler. The mixture is cured to encase the fiber mats in a porous polymer matrix. The resulting composite intermediate product is heated to drive off the pore-forming agent and pyrolyze the matrix to a porous carbon char.

In the unmodified version of the process, the next step is to infiltrate the porous carbon matrix with liquid silicon or a silicon alloy. The liquid silicon or silicon alloy reacts with the carbon matrix to form a matrix of mostly silicon carbide with other phases mixed in. Unfortunately, the liquid silicon also reacts with the carbon-rich surface layers on the SiC fibers, consuming those

layers and attacking the underlying SiC fiber material. Such attack can be prevented by use of protective coats on the SiC fibers, but of course, protective coating increases the cost.

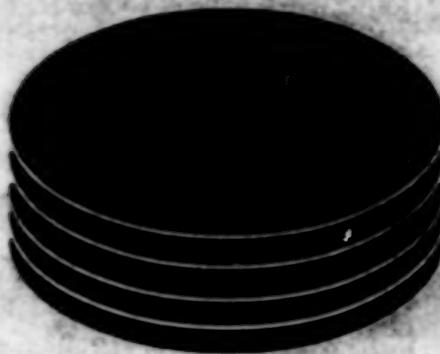
In the modified version of the process, the damage to the SiC fibers is reduced, without need for protective coating. The modification consists of the addition of a second polymer-infiltration-and-pyrolysis step to increase the carbon content of the porous matrix before infiltration with liquid silicon or silicon alloy. The polymer-forming mixture used in this step is similar to that used in the first step except that it does not contain carbon powder. Next,

the matrix is infiltrated with liquid silicon or silicon alloy. As before, the silicon reacts with the carbon matrix material to form a silicon carbide matrix, but in this case, the silicon carbide fibers and their carbon-rich surface layers remain intact.

This work was done by Donald R. Behrendt of Lewis Research Center and Mrityunjay Singh of NYMA, Inc. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Lewis Research Center [see page 1]. Refer to LEW-15690.

BLANK PAGE



Computer Programs

Electronic Components and Circuits

41 Program Calculates Current Densities of Electronic Designs

Physical Sciences

41 Program Models a Laser Beam Focused in an Aerosol Spray
42 TSS — Thermal Synthesizer System

Materials

43 Program for Evaluation of Reliability of Ceramic Parts

Mechanics

44 Computing Operating Characteristics of Bearing/Shafit Systems

Mathematics and Information Sciences

44 Computing Displacements and Strains From Video Images

BLANK PAGE

Computer Programs

These programs may be obtained from COSMIC. Please contact

COSMIC

Computer Services Annex
University of Georgia
Athens, GA 30602
Telephone No. (404) 542-3265.

Electronic Components and Circuits

Program Calculates Current Densities of Electronic Designs

The current in each component is averaged over the area occupied by the component.

The PDENSITY computer program calculates current densities for use in calculating power densities of electronic designs. PDENSITY reads a parts-list file for a given design, a file that contains the current required for each part, and a file that contains the size of each part. Then for each part in the design, the program calculates the current density in units of milliamperes per square inch (1 mA/in.² = 0.155 mA/cm²).

The user invokes PDENSITY with the input parts list, the current-requirements file, the size file, and the name of the output file. Any parts not in the current-requirements file are reported, and the current-requirements file can be modified for particular needs.

PDENSITY is written by use of the AWK utility for Sun4-series computers running SunOS 4.x and IBM PC-series and compatible computers running MS-DOS. The program was written to be used with the OrCed Schematic Capture Program, but also can be used with a list of parts in ASCII format. The UNIX "hawk" utility is used to run the Sun version of PDENSITY. The Sun version of the program (NPO-19588) has been successfully implemented on a Sun4-series computer running SunOS 4.1.3. The PC version of the program (NPO-19171) has been successfully implemented on a DECpc 486-series computer running MS-DOS v6.2. A sample executable file is provided for the PC version of PDENSITY. The standard distribution medium for the Sun version of PDENSITY is a 3.5-in. (8.89-cm), 1.44MB diskette in UNIX tar format. The standard distribution medium for the PC version of PDENSITY is one 3.5-in. (8.89-cm), 1.44MB, MS-DOS format diskette. An electronic copy of the documentation in Microsoft Word v2.0 for MS-DOS format is included on all distribution media. PDENSITY was released in 1995 and is a copyrighted work with all copyright vested in NASA.

This program was written by Brian Cox of Caltech for NASA's Jet Propulsion Laboratory.

For further information on NPO-19588, see TSP's [page 1].

For further information on NPO-19171, see TSP's [page 1].
NPO-19588/NPO-19171

Physical Sciences

Program Models a Laser Beam Focused in an Aerosol Spray

A Monte Carlo analysis is performed on packets of light.

The Program for the Analysis of a Laser Beam Focused Within an Aerosol Spray (FLSPRY) was developed for theoretical analysis of the propagation of a laser pulse optically focused within an aerosol spray. FLSPRY can be applied, for example, to analyze a laser ignition arrangement in which a focused laser pulse would be used to ignite a liquid aerosol fuel spray. Scattering and absorption of laser light by the individual aerosol droplets are evaluated by use of electromagnetic Lorenz-Mie theory.

Initially, propagation of a laser beam is modeled by use of a simple modified parallel-ray theory, which would limit applications to relatively "optically thin" sprays. To keep from limiting the application of the program to "optically thin" sprays, a Monte Carlo analysis of the laser beam has been adapted to FLSPRY. The Monte Carlo analysis follows individual "light packets" through a dense spray of defined thickness and composition.

The spray is assumed to be monodisperse and homogeneous, and the light packets enter along the same axis with

perpendicular incidence. The Monte Carlo algorithm was modified to account for the spatial variation of intensity and the variable angle of incidence at the illuminated side of the spray. An appropriately weighted Monte Carlo decision is used to determine the position of incidence for each light packet. After a large number of attempts, a Gaussian intensity profile is generated.

Inputs to FLSPRY include the volume fraction of liquid, average size of droplets, droplet-size distribution, laser wavelength, laser-pulse energy, laser-pulse duration, focal length of the lens, diameter of the beam incident on the lens, and the choice of aerosol liquid and surrounding gaseous medium. Other inputs include Monte Carlo analysis indicator, radius (cm) for Monte Carlo analysis, location of Monte Carlo analysis (cm) relative to the lens, number of Monte Carlo attempts, and random-number-generator seed for Monte Carlo analysis. Also necessary for Monte Carlo analysis are the x_{max} , x_{min} , y_{max} , and y_{min} thicknesses of the spray (cm). The output of the computer program includes, as a function of spatial position along the laser propagation axis within the spray, the laser-pulse intensity and energy, the overall volumetric absorption of laser energy by the aerosol liquid and by the gaseous medium, and the overall average temperature rise of the aerosol liquid and of the gaseous medium.

FLSPRY is written in FORTRAN 77 for both UNIX-based computers and DEC VAX-series computers. FLSPRY requires a compiler that supports NAMELIST input. The VAX version of the program (LEW-16051) has been successfully implemented on DEC VAX-series computers running VMS v5.5. The UNIX version of the program (LEW-16085) has been successfully implemented on SGI IRIS-series computers running IRIX 5.2, and DEC ALPHA AXP-series computers running OSF/1. The author of the program has also successfully implemented FLSPRY on IBM RS/6000-series computers running AIX v3.2, using the FORTRAN 77 v2.03 compiler. The standard distribution medium for the VAX version of FLSPRY is a 1,600-bit/in. (830-bit/cm), 9-track magnetic tape in DEC VAX BACKUP format. It is also available on a TK50 magnetic-tape cartridge in DEC VAX BACKUP format. The standard distribution medium for the UNIX version of FLSPRY is a 0.25-in. (6.35-mm) streaming-magnetic-tape cartridge (Sun

QIC-24) in UNIX tar format. Upon request, the UNIX version of FLSPRY can be provided in UNIX tar format on alternate media. FLSPRY was released in 1994.

This program was written by J. P. Barton of the University of Nebraska in Lincoln for Lewis Research Center.

For further information on LEW-16051, see TSP's [page 1].

For further information on LEW-16065, see TSP's [page 1].
LEW-16051/65

TSS — Thermal Synthesizer System

This software can greatly expedite the thermal-engineering process.

The Thermal Synthesizer System (TSS) is an integrated set of thermal-analysis application programs that is designed to solve problems encountered by thermal engineers. TSS combines the functionality of Systems Improved Numerical Differencing Analyzer/Fluid Integrator (SINDA/FLUINT) and radiation analysis with a friendly and easily understood user-interface environment coupled with powerful interactive color graphics and geometric modeling capability. The effort expended on developing this means of interaction has resulted in a system that is easy and straightforward to use. This system will enable thermal engineers to spend more of their time solving engineering problems instead of laboriously constructing and verifying math models.

TSS functionality is logically grouped into a number of application programs. The Executive application program serves as the starting point for all other TSS application programs. This application program also encapsulates UNIX functionality into a user-friendly interface. The Geometry application program facilitates construction, verification, and display of thermal-radiation mathematical models as well as solid and shell models for subsequent processing by the Radiation Conductance (RADK), Heatrate, Conductance-Capacitance, and Image application programs.

The RADK application program allows calculation of thermal-radiation form factors, interchange factors, and radiation conductances through use of Monte Carlo ray-tracing techniques. RADK output is compatible with SINDA/FLUINT, also included in TSS. The Orbit application program provides a means of creating and visualizing orbits for subsequent

use by the Animation and Heatrate application programs. The Orbit application program operates in six different modes: basic, Keplerian, Sun-referenced, date-dependent, heliocentric, and planetary surface. The Animation application program facilitates the articulation of model geometries (in either an orbiting or a geometry-only mode) for subsequent processing by the RADK and Heatrate application programs. Each object within a geometric model can be articulated in one of four animation modes: velocity plus rotation rate, positional interpolation, inertial tracking, or point tracking.

Users may specify artificial radiant-heat sources within the Heatsource application program. Four different types of artificial sources may be specified: infinite parallel, point, nodal diffuse, and infinite diffuse. Both solar and infrared spectrum sources are supported. The Heatrate application program enables characterization of orbit environments, planetary surface environments, and artificial heating. Orbit-environment calculations provide for computation of solar, albedo, and planetary heating components. The Conductance-Capacitance application program facilitates the creation of a SINDA/FLUINT-compatible thermal-network model as derived from the geometry specified by the user. Nodal capacitances and intraobject conductances are automatically calculated. Interobject conduction paths are easily generated.

The SINDA/FLUINT application program, included with TSS, is the same SINDA/FLUINT program that is available on many platforms (Version 2.6). Models created within TSS can be processed within SINDA/FLUINT or exported to external systems. Conversely, any existing SINDA/FLUINT model can be ported into the TSS system and processed. The FLUINT Modeler application program enables the construction, display, and verification of single- or two-phase flow networks and creates SINDA/FLUINT-compatible input decks. This application program also supports postprocessing of the model results onto the geometric representation of the flow network.

The Thermal Radiation Analysis System (TRASYS) to Geometry Translator converts models in the TRASYS input-deck format into the TSS Geometry Input Language. The Geometry to TRASYS Translator converts models in the TSS Geometry Input Language format into the TRASYS input-deck format.

Results from TSS analyses, ASCII data, and flight data can be plotted in

the XY Plot application program, which is included with TSS. Users can work with up to eight different sets of data simultaneously. Customization of plots, comparison-plotting of sets of data, and curve fitting are also available. Geometries defined by users can be processed in the Image application program to provide photorealistic images to support engineering analysis, presentations, and documentation.

TSS has been under development for many years. It was first released in 1993. TSS is available from COSMIC (v2.0) by license for a period of ten years to approved licensees. The licensed program product includes the executable code and one copy of the supporting documentation. Additional copies of the documentation may be purchased separately at any time.

TSS is written in FORTRAN and C language for Apollo DN10000VS and HP 700-series computers running Domain and UNIX operating systems, respectively. The program requires 32MB of random-access memory and at least 500MB of disk space. For the Apollo workstations, TSS requires 40-bit-plane graphics, Domain 3-D graphics Metatile and Domain Dialogue. For the HP-series workstations, TSS requires the CRX 24Z graphics upgrade, HPUX 9.05, and the PHIGS runtime library. A FORTRAN-77 compiler is required for the SINDA/FLUINT application.

The standard distribution medium is magnetic tape. For the Apollo computers, the program is distributed on 60MB, 0.25-in. (6.35-mm) streaming-magnetic-tape cartridge in rbae format, which is readable on a Domain Dialogue 60MB tape drive. For the HP-series workstations, the program is distributed on 4-mm tape in tar format, readable under standard UNIX environments.

This program was developed by Edward Chimenti, Steven Rickman, Robert Vogt, and Carlos R. Ortiz Longo of Johnson Space Center; Noel Bauman, Joseph Lepore, Phil Mackey, James Pavlovsky II, Mark Welch, and Peter Fogerson of Lockheed Engineering and Sciences Co.; and Mark Dawber, Cynthia Jane Fong, Peter Hecke, Susan Morrison, Ernie Castillo, Zu Chou, Lawrence Fried, Jerry Howard, Mike Lombardi, Jack Middleton, Tim Panczak, and Brett Preston of Lockheed Missiles and Space Co. For further information, see TSP's [page 1].

MSC-22420

Materials

Program for Evaluation of Reliability of Ceramic Parts

This program predicts the probability of failure as a function of time in service.

Successful application of advanced ceramics depends on proper characterization of material properties and the use of probabilistic methodology for designing parts made of brittle materials. CARES/LIFE is a computer program for predicting the probability of failure of a monolithic ceramic component as a function of its time in service.

Probabilistic component design involves predicting the probability of failure of a thermomechanically loaded component from specimen rupture data. Typically, experiments to acquire these data are performed by use of many simple-geometry flexural or tensile test specimens. A static, dynamic, or cyclic load is applied to each specimen until fracture. Statistical strength fatigue parameters are then determined from these data.

Using these parameters and the results obtained from a finite-element analysis, the time-dependent reliability for a complex component geometry and loading is then predicted. Appropriate design changes are made until an acceptable probability of failure has been reached. This design methodology combines the statistical nature of strength-controlling flaws with the mechanics of crack growth to allow for multiaxial stress states, concurrent (simultaneously occurring) flaw populations, and subcritical crack growth (SCG).

CARES/LIFE predicts the probability of failure of a monolithic ceramic component as a function of service time. It assesses the risk that the component will fracture prematurely as a result of SCG. The effect of proof testing of components prior to service is also considered. CARES/LIFE is coupled to such commercially available finite-element programs such as ANSYS, ABAQUS, MARC, MSC/NASTRAN, and COSMOS/M. It also retains all of the capabilities of the previous CARES code, which include estimation of fast-fracture com-

ponent reliability and Weibull parameters from inert strength (without SCG contributing to failure) specimen data. CARES/LIFE can estimate parameters that characterize SCG from specimen data as well.

Finite-element heat transfer and linear-elastic stress analyses are used to determine the distributions of temperature and stress in a component. The reliability at each element is calculated under the assumption that randomly distributed volume flaws and/or surface flaws control the failure response. The overall component reliability is the product of all the element-survival probabilities. CARES/LIFE generates a data file containing element risk-of-rupture intensities (a local measure of reliability) for graphical rendering of the critical regions of a structure.

CARES/LIFE describes the probabilistic nature of the strength of a material, using the Weibull cumulative distribution function. The Weibull equation is based on the weakest-link theory (WLT), in which it is assumed that the structure is analogous to a chain with many links. Each link may have a different limiting strength. When a load is applied to the structure such that the weakest link fails, then the structure fails.

The effect of multiaxial stresses on reliability is predicted by using the principle of independent action (PIA), the Weibull normal-stress-averaging method (NSA), or the Batdorf theory. For the PIA model, the reliability of a component under multiaxial stresses is the product of the reliability of the individual principal stresses acting independently. The NSA method involves the integration and averaging of tensile normal stress components evaluated about all possible orientations and locations. This approach is a special case of the more general Batdorf theory and involves the assumption that the material is insensitive to shear.

The Batdorf theory combines the WLT and linear elastic fracture mechanics (LEFM). A flaw geometry selected by the user and a mixed-mode fracture criterion are required to model volume- or surface-strength-limiting defects. "Mixed-mode fracture" refers to the ability of a crack to grow under the combined actions of a normal load (opening mode) and shear load (sliding and tearing modes) on the crack face.

In CARES/LIFE, the relations that describe subcritical crack growth are directly incorporated into the PIA, NSA,

and Batdorf theories. Subcritical crack growth is modeled with the power law, the Paris law, and the Walker law for static and constant-amplitude cyclic loading. These laws use experimentally determined parameters that are material- and environment-sensitive.

Predicted lifetime reliability of structural ceramic components depends on Weibull and fatigue parameters estimated from widely used tests that involve flexural or tensile specimens. CARES/LIFE estimates fatigue parameters from naturally flawed specimens ruptured under static, cyclic, or dynamic (constant stress rate) loading. Fatigue and Weibull parameters are calculated from rupture data on three-point and four-point flexure bars as well as tensile specimens. For other specimen geometries, finite-element models of the specimens are also required when estimating these parameters.

CARES/LIFE is written in ANSI FORTRAN 77 to be machine-independent. Therefore, the program will run on any computer in which sufficient addressable memory (at least 8MB) and a FORTRAN 77 compiler are available. The CARES/LIFE software package has been successfully implemented on DEC VAX computers running VMS, Sun4-series computers running SunOS 4.1.3, Silicon Graphics computers running IRIX 5.2, and DECstation-series computers running DEC RISC ULTRIX. Sample output files from the ANSYS (Swanson Analysis Systems, Inc.), MSC/NASTRAN (MacNeal-Schwendler Corp.), and ABAQUS (Hibbit, Karlsson, and Sorenson, Inc.) commercially available finite-element-analysis software packages are included with the distribution. For an IBM-compatible personal computer with a minimum 640K memory, a limited program is available (CARES/PC, COSMIC number LEW-15248). The standard distribution medium for CARES/LIFE is a 0.25-in. (6.35-mm) streaming-magnetic-tape cartridge (Sun QIC-24) in UNIX tar format. Alternate distribution media and formats are available upon request. CARES/LIFE was released in 1994.

This program was written by N. Nemeth, L. A. Janosik, and J. P. Gyekenyesi of Lewis Research Center and Lynn M. Powers of Cleveland State University. For further information, see TSP's [page 1].
LEW-16018

Mechanics

Computing Operating Characteristics of Bearing/Shaft Systems

SHABERTH incorporates up-to-date mathematical models of thermal and mechanical phenomena.

The SHABERTH computer program predicts operating characteristics of bearings in a multibearing load-support system. Lubricated and nonlubricated bearings can be modeled. SHABERTH calculates the loads, torques, temperatures, and fatigue lives of ball and/or roller bearings on a single shaft. The program also provides for analysis of the reaction of the system to the termination of the supply of lubricant to the bearings and other lubricated mechanical elements. SHABERTH has proven to be valuable in the design and analysis of shaft/bearing systems.

The SHABERTH program is structured with four nested calculation schemes. A thermal scheme performs both steady-state and transient calculations that predict temperatures in the system in a given operating state. A bearing-dimensional-equilibrium scheme uses (1) bearing temperatures predicted by the temperature-mapping subprograms in the thermal scheme and (2) distributions of loads on rolling-element raceways as predicted by a bearing subprogram, to calculate (3) diametral clearances of bearings in a given operating state. A shaft/bearing-system load-equilibrium scheme calculates positions of inner rings of bearings relative to the respective outer rings, such that the external loads applied to the shaft are brought into equilibrium by the rolling-element loads that develop at each inner ring in a given operating state. A bearing-rolling-element-and-cage load-equilibrium scheme calculates the equilibrium positions and the speeds of rotation of rolling elements and cages on the basis of relative positions of inner and outer rings, inertial effects, and friction.

The current version of SHABERTH contains ball-bearing subprograms that are enhanced, in comparison with similar programs, by the incorporation of several mathematical models. These include

- An elastohydrodynamic (EHD) film-thickness model that accounts for

heating and lubricant-film starvation in the contact area;

- A new model for traction combined with an asperity-load-sharing model;
- A model for the hydrodynamic rolling and shear forces in the inlet zones of lubricated contacts, which model accounts for the degree of lubricant-film starvation;
- A model of normal and friction forces between a ball and a cage pocket, which model accounts for the transition between the hydrodynamic and elastohydrodynamic regimes of lubrication; and
- A model of the effect, on fatigue life, of the ratio between the thickness of the EHD plateau film and the composite surface roughness.

SHABERTH is intended to be as general as possible. The models in SHABERTH provide for the complete mathematical simulation of real physical systems. A simulated system is limited to a maximum of five bearings supporting a shaft, a maximum of 30 rolling elements per bearing, and a maximum of 100 temperature nodes. The structure of the SHABERTH program is modular and has been designed to enable refinement and replacement of various component models as needs and opportunities develop.

SHABERTH is available in Cray and IBM PC versions. A preprocessor subprogram is included in the IBM PC version to provide a user-friendly means of developing SHABERTH models and executing the resulting code. The preprocessor enables the user to create and modify data files with minimal effort and a reduced chance for errors. Data are utilized as they are entered; the preprocessor then decides what additional data are needed to complete the model in question. Only these needed data are requested. The preprocessor can accommodate input of data for any shaft/bearing-system mathematical model compatible with SHABERTH. The system to be simulated can include ball bearings, cylindrical roller bearings, and/or tapered roller bearings.

The two versions of SHABERTH are available from COSMIC. The Cray version [LEW-14960, "Computing Thermal Performances of Shafts and Bearings," *NASA Tech Briefs*, Vol. 16, No. 3 (March, 1992), page 67] requires a random-access memory (RAM) of 176K of 64-bit words. The IBM PC version (MFS-28818, this article) is written for IBM PC-series and compatible computers running MS-DOS and includes a sample MS-DOS executable code. For execution, the PC

version requires at least 1MB of RAM and computer that contains an 80386 or 486 processor with an 80 x 87 math coprocessor. The standard distribution medium for the IBM PC version is a set of two 5.25-in. (13.335-cm), 360K, MS-DOS-format diskettes. The contents of the diskettes are compressed by use of the PKWARE archiving software tools. The utility software to unarchive the files, PKUNZIP.EXE, is included. The standard distribution medium for the Cray version is also a 5.25-in. (13.335-cm), 360K, MS-DOS-format diskette, but alternate distribution media and formats are available upon request.

The original version of SHABERTH was developed in FORTRAN IV at Lewis Research Center for use on a UNIVAC 1100-series computer. The Cray version was released in 1988, and was updated in 1990 to incorporate fluid-rheological data for Rocket Propellant 1 (RP-1), thereby enabling the analysis of bearings lubricated with RP-1. The PC version is a port of the 1990 Cray version and was developed in 1992 by SRS Technologies under contract to NASA's Marshall Space Flight Center.

This program was written by James D. Moore of Marshall Space Flight Center. For further information, see TSP's [page 1].

MFS-28818

Mathematics and Information Sciences

Computing Displacements and Strains From Video Images

There is no need for custom or expensive equipment.

Subpixel digital video image correlation (SDVIC) is a technique for measuring in-plane displacements on the surfaces of objects under loads, without contact. This technique can be used for analyses of experimental research specimens or actual service structures of virtually any size or material. Only minimal preparation of test objects is needed, and there is no need to isolate test objects from minor vibrations or fluctuating temperatures. The technique is implemented by the SDVIC software,

which produces color-graduated, full-field representations of in-plane displacements and their partial derivatives with respect to position along both principal directions in each image plane. From these representations, linear strains, shear strains, and rotation fields can be determined.

The SDVIC software is based upon an algorithm for locating small regions of a random pattern after the pattern has been deformed. The SDVIC software determines values for in-plane displacements and strains by correlating the positions of subsets of pixels in the original image with those in the deformed image on the basis of pixel gray levels in the digitized versions of the images. The sizes of subsets to be pattern-matched can be customized, thus enabling the correlation of many random patterns. Several techniques, including a coarse-fine search and the Newton-Raphson method, are available for subset pattern matching. A bilinear interpolation routine provides subpixel resolution necessary for most displacement and strain measurements.

To utilize SDVIC, one needs a black-and-white video camera to acquire images of the test article before and after loading. Camera-lens combinations can

be altered to enable analyses of test articles of virtually any size. A personal-computer-based video circuit board is typically used to digitize the images. The images are then correlated by a personal computer equipped with the SDVIC software. All of the components needed for a typical system are commercially available. There is no need for custom equipment, which can be expensive or difficult to find.

In order to correlate properly images based upon subset gray-level patterns, one must have a characteristic pattern on the surface of the test article. This pattern can be one that occurs naturally, (for example, the microscopic grains and grain boundaries of a metal as viewed under magnification), or the pattern can be applied by dusting the specimen with fine particles or spraying the specimen with white and black spray paint. Application of a pattern typically requires only a few minutes and a few dollars for several specimens.

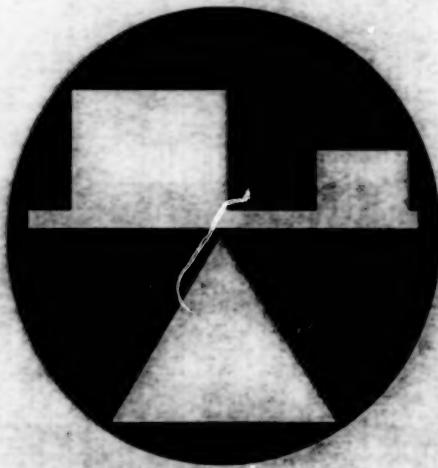
SDVIC is written in C language for IBM PC-series and compatible computers running MS-DOS. The minimum requirements for execution of SDVIC include an 80386 processor, Microsoft Windows v3.1, an SVGA monitor, 2MB of random-

access memory (RAM), and 5MB of swap space. For best results, an 80486/66-MHz processor, 8MB of RAM, and 32MB of swap space are highly recommended. SDVIC can be executed with Windows v3.1 in the enhanced mode. A sample executable code is provided on the distribution medium. If there is need to recompile the source code, then Microsoft Software Developers Kit v3.1 and Microsoft C compiler v6.0 are required. An electronic copy of the documentation in Microsoft Word for Windows v2.0b format is included on the distribution medium. The standard distribution medium for SDVIC is a set of four 3.5-in. (8.89-cm), 1.44MB, MS-DOS-format diskettes. This version of SDVIC was developed in 1983.

This program was written by Stephen R. McNeill of the University of South Carolina, Samuel S. Russell of Marshall Space Flight Center, and Matthew D. Lansing of the University of Alabama in Huntsville. For further information, see TSP's [page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-26289.

BLANK PAGE



Mechanics

Hardware, Techniques, and Processes

- 49 Refinement of Hexahedral Cells in Euler Flow Computations
- 50 Sensing Position With Approximately Constant Contact Force
- 50 Multigrid, Fractional-Step Computation of Flow
- 51 Bearings Incorporating Deadband Rollers

Books and Reports

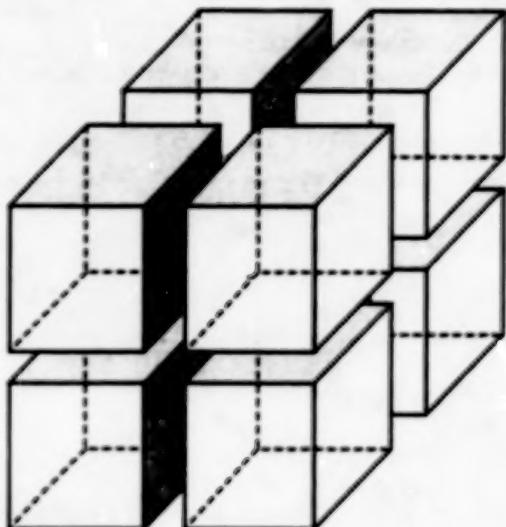
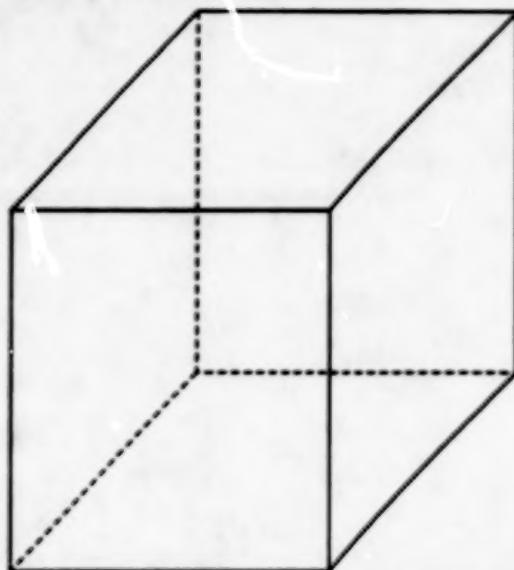
- 51 Non-Newtonian Effects in Viscous Flows
- 52 Theory of Controlling Spacecraft Motion With Pulsed Thrusters

BLANK PAGE

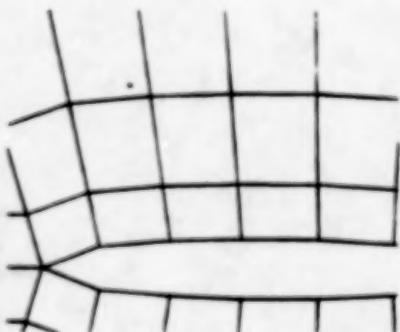
Refinement of Hexahedral Cells in Euler Flow Computations

Computational grids are refined where necessary to resolve details more accurately.

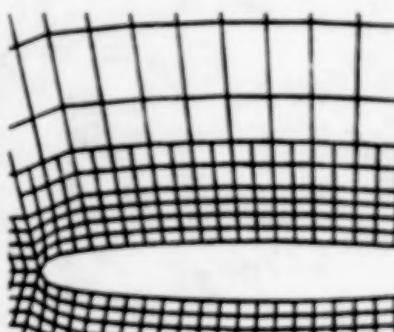
Ames Research Center,
Moffett Field, California



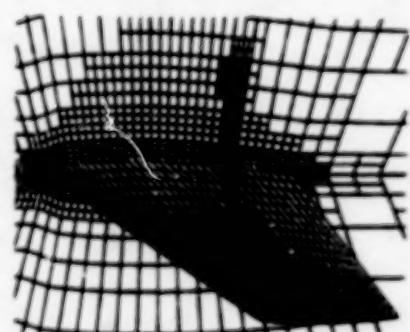
SUBDIVISION OF A HEXAHEDRAL CELL



Original Grid at the
50-Percent-of-Span Plane



Grid at the 50-Percent-of-Span
Plane After the Second Step
of Refinement



Part of the Final Refined Grid

ASPECTS OF REFINEMENT OF THE COMPUTATIONAL GRID ABOUT AN AIRPLANE WING

Hexahedral Cells are subdivided, each into eight smaller cells, as needed to refine the computational grid in regions of high flow gradients.

The Topologically Independent Grid, Euler Refinement (TIGER) computer program solves the Euler equations of three-dimensional, unsteady flow of an inviscid, compressible fluid by numerical integration on an unstructured hexahedral coordinate grid that can be refined where necessary to resolve shocks and other details. (As used here, "unstructured" means arbitrarily and/or complicatedly shaped, usually to fit an irregularly shaped body in the flow field.) The Grid Interactive Refinement and Flow-Field Examination (GIRAFFE) computer program was written in conjunction with the TIGER program to display the computed flow-field data and to assist the

researcher in verifying specified boundary conditions and refining the grid.

The TIGER program is based on Jameson's FL057 program, which solves the Euler equations on a structured grid by use of a finite-volume, four-stage Runge-Kutta numerical-integration algorithm. The development of TIGER involved modifications to incorporate conservative calculations of mass, momentum, and energy fluxes across the faces of multiply connected grid cells. The unstructured-grid nature of TIGER makes it amenable to local refinement of the grid, as specified by the researcher.

Each grid cell can be split into eight smaller cells (see figure). If further refine-

ment is needed, the smaller cells can be subdivided similarly. Refinement of the grid appropriate to each flow-field condition and tailored to capture shocks or other specific flow features can be based on a variety of criteria; for example, the ratio between the local and free-stream total pressures, the difference between the maximum and minimum pressures in the vicinity, or the preliminary calculated location of a shock.

The TIGER program takes about the same amount of time to compute a flow as do conventional flow-simulating programs that solve the Euler and Navier-Stokes equations on structured (that is, regular) grids. The TIGER program

involves the use of simple, straightforward programming techniques to take advantage of the vector-processing capabilities of the CRAY Y-MP computer. The TIGER program requires approximately five times more memory per grid point than a conventional structured-grid program does, but in TIGER, it is not necessary to refine the grid throughout the entire flow

domain. When the refinement of the grid is based on flow parameters specified by the researcher, points of the resulting grid are usually distributed more appropriately for the problem, yielding a net decrease in the total memory required.

This work was done by John E. Melton and Gelsomina Cappuccio of Ames Research Center and Scott D. Thomas

of Sterling Federal Systems. Further information may be found in AIAA paper 91-0637: "Unstructured Euler Flow Solutions Using Hexahedral Cell Refinement."

Copies may be purchased [prepayment required] from AEROPLUS, Burlingame, California 94010, Telephone No. (800) 662-2376, Fax No. (415) 259-5047. ARC-13077

Sensing Position With Approximately Constant Contact Force

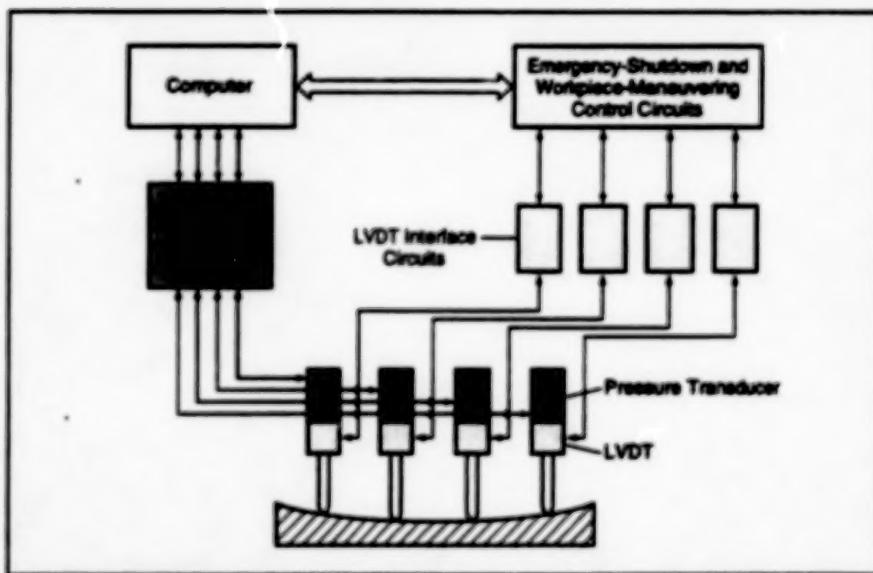
LVDTs are held lightly against the surface to be measured.

A computer-controlled electromechanical system uses a number of linear variable-differential transformers (LVDTs) to measure the axial positions of selected points on the surface of a lens, mirror, or other precise optical component with a high finish. The system maintains the contact force of each probe at a level between the weights of 9 and 13 grams; this level of force is just large enough to ensure consistent measurements but not so large as to cause damage to the surface when the probe is compressed or extended during axial positioning of the optic, over an axial range of 7 mm, between measurement positions in contact with the optic surface when the optic is rotated between measurement positions.

The contact tip of each LVDT is a 4-mm ruby ball. Each LVDT is spring-loaded in retraction from the surface to be measured. The probe is driven pneumatically, against the spring loading, into contact with the surface. The pneumatic pressure applied to each probe is sensed and adjusted via an electromechanical pressure transducer. Operating under computer control, the transducers for all the probes maintain the pressures that produce the required contact forces.

The output of each LVDT is a voltage indicative of the position of its tip. During the initial approach to the surface to be measured, the probes are retracted. The pneumatic pressure driving each probe is increased in steps to extend the probe in steps of 0.006 in. (-0.15 mm).

Marshall Space Flight Center, Alabama



The pressure applied to pneumatically driven LVDTs are adjusted under computer control to maintain small, approximately constant contact forces as the positions of the LVDT tips vary.

Eventually, the position reading from a probe stops increasing, signifying that the probe has made contact with the surface. Thereafter, the computer maintains the pressure at the level that corresponds to the required contact force. In computing this pressure, the computer accounts for the increase in spring retraction force with increasing extension of the tip or decrease in spring retraction force with compression of the tip.

This work was done by Jay Sturdevant of Hughes Aircraft Co. for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Title to this invention has been waived under the provisions of the National Aeronautics and Space Act (42 U.S.C. 2457(f)) to Hughes Aircraft Co. Inquiries concerning licenses for its commercial development should be addressed to

Hughes Aircraft Co.
Attn: Jay Sturdevant
7200 Hughes Terrace
P.O. Box 45066
Los Angeles, CA 90045-0066
Refer to MFS-28936, volume and number of this NASA Tech Briefs issue, and the page number.

Multigrid, Fractional-Step Computation of Flow

Computation of incompressible flow to high resolution takes less time.

The speed of a computer code that solves the three-dimensional Navier-Stokes equations of flow of an incom-

pressible fluid by the fractional-step method has been increased significantly by use of multigrid procedures. In the

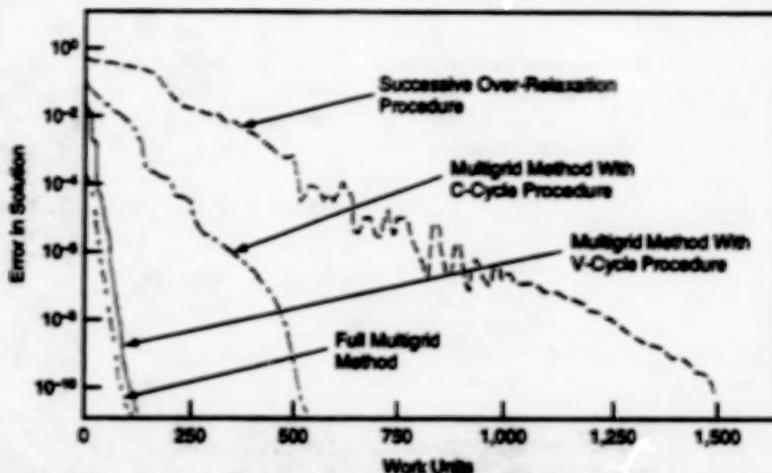
Ames Research Center, Moffett Field, California

specific version of the fractional-step method, the equations are solved on a general nonorthogonal curvilinear coor-

dinate grid, using volume fluxes (instead of velocities). At each time step, computations are performed in two fractional steps. In the first fractional step, the equations of conservation of momentum are solved by use of the gradient of pressure from the previous time step via an explicit approximate-factorization method, yielding an approximate flow field that does not satisfy the equation of conservation of mass. In the second fractional step, a discrete Poisson-like equation with Neumann-type boundary conditions, formed by combining the equations of conservation of momentum and mass, is solved iteratively.

The number of arithmetic operations of, and therefore the time consumed by, the iterative solution of the Poisson-like equation increases with the number of grid points; heretofore, it has usually amounted to more than 80 percent of the total time of a fractional-step computation. The multigrid method is the only one that provides solution procedures that both are applicable to general curvilinear coordinates and entail numbers of arithmetic operations that increase no more than linearly with the number of grid points. Furthermore, multigrid techniques for the solution of elliptic types of equations similar to the Poisson-like equation are well developed.

To test the ability of the multigrid method to accelerate convergence toward the solution of the Poisson-like equation, several incompressible flows were computed. These included two-dimensional flow about an elliptical airfoil, two-dimensional flow between two eccentric cylinders, and three-dimensional flow in a square duct with a bend. In each case, the computational time need-



This History of Convergence toward the solution of the Poisson-like equation was obtained in a computation of flow between eccentric cylinders on a grid of 256×256 points, using four different solution procedures. A work unit (the scale of the abscissa) is a measure of computational effort, essentially proportional to the number of arithmetic operations.

ed to solve the Poisson-like equation was reduced by an order of magnitude (see figure), while the overall computation time was reduced by a factor of 3 to 4.

The solution of the Poisson-like equation by use of the multigrid method consumed less than 25 percent of the total central-processing-unit time. The computational work was found to be approximately proportional to N (where N is the number of grid points), while the total central-processing-unit time on a Cray Y-MP vector computer was found to be approximately proportional to $N^{0.75}$. Thus, it appears that the combination of multigrid and fractional-step methods provides a viable alternative for the computation of complicated flow fields with very large numbers of grid points.

This work was done by Dochan Kwak of Ames Research Center and Moshe Rosenfeld of Tel Aviv University. Further information may be found in AIAA paper 92-0185, "Multigrid Acceleration of a Fractional-Step Solver in Generalized Curvilinear Coordinate Systems."

Copies may be purchased [prepayment required] from AEROPLUS, Burlingame, California 94010, Telephone No. (800) 652-2376, Fax No. (415) 259-5047.

This is the invention of a NASA employee, and a patent application has been filed. Inquiries concerning license for its commercial development may be addressed to the inventor:

Dochan Kwak
4153 Georgia Ave.
Palo Alto, CA 94306
Refer to ARC-13196.

Bearings Incorporating Deadband Rollers

Bearings in a high-pressure turbopump have been redesigned to incorporate rollers that allow limited axial motion within a small deadband. The design does not permit radial deadband motion, which is undesired. The axial deadband motion would be used for rotor-thrust-balance control. The axial-deadbond rollers would

be supported by a multisided bearing carrier. This design would eliminate some nonlinearities in the dynamics of the pump rotor and assist in suppressing vibrations at harmonics of the frequency of rotation.

This work was done by Guy V. Gagliardi of United Technologies for Marshall Space Flight Center. To obtain a copy of

the report, "Bearing Deadband Rollers," see TSP's [page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-28914.

Books and Reports

Non-Newtonian Effects in Viscous Flows

A report presents a theoretical study that addresses the persistent problem of

explaining the random aspects of the flows of real fluids in terms of the classical governing equations. A review of the classical mathematical formalism of fluid dynamics leads to the conclusion that some of the

physically unrealistic aspects of classically computed flows (for example, infinite times of approaching equilibria, and fully deterministic solutions of the Navier-Stokes equations) can be removed by relaxing the

Lipschitz condition, which is a requirement that the derivatives of the solutions of the differential equations of flow be bounded. (The Lipschitz condition guarantees the uniqueness of the solution under fixed initial conditions, whereas what is sought here is the nonuniqueness characteristic of real, partly random flows.) The relaxation of the Lipschitz condition can be interpreted in physical terms as the incorporation of an infinitesimal static friction in the constitutive law of a viscous fluid. A non-Lipschitzian modified version of the Navier-Stokes equations is introduced and shown to yield nonunique solutions; in other words, solutions with the desired fundamental new property of stochasticity. All the new stochastic effects emerge within vanishingly small neighborhoods of equilibrium dynamical states; these infinitesimal neighborhoods are the only domains where the modified governing equations differ from the classical. This means that the formal differences between the solutions to the classical and modified equations may be not detectable in domains that do not include equilibrium states.

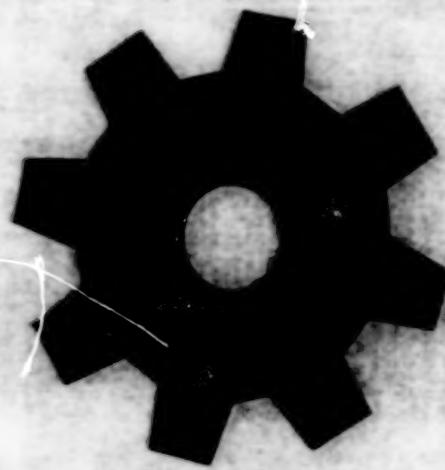
This work was done by Michail Zak of Caltech and Ronald E. Meyers of the U. S. Army Research Laboratory for NASA's Jet Propulsion Laboratory. To obtain a copy of the report, "Non-Newtonian Effects in Viscous Flows," see TSP's [page 1].
NPO-19728

Theory of Controlling Spacecraft Motion With Pulsed Thrusters

A report presents a new class of flight-control laws for making a spacecraft follow a desired trajectory by use of pulsed thrusters during such maneuvers as automated rendezvous on orbit and soft landing on a planet. The guidance problem is solved with an extension of robust nonlinear control theory, resulting in algorithms that determine the duration and frequency of thruster firings from estimates of spacecraft position and velocity. These algorithms incorporate a complete nonlinear dynamical model of the spacecraft motion, and enable analytical characterization of transient-error behavior, limit-

cycle deadband region, and the domain of possible terminal states. In addition, it is shown that the desired performance is achievable in the presence of dynamical-modeling errors with known bounds, and that the effects of navigation and attitude-sensing errors can be minimized, given that appropriate bounds can be determined for these quantities. Although the emphasis in the report is on guidance, it is ultimately shown that the same techniques are also useful for attitude control and station keeping. A realistic application is illustrated via computer simulation of a soft landing, on Mars, of a robotic spacecraft equipped with an inertial-guidance system.

This work was done by Sam W. Thurman of Caltech and Henryk Raszner of the University of Southern California for NASA's Jet Propulsion Laboratory. To obtain a copy of the report, "A New Pulse-Modulation Technique for Guidance and Control of Automated Spacecraft," see TSP's [page 1].
NPO-19695



Machinery

Hardware, Techniques, and Processes

- 55 Helium-Recycling Plant
- 56 Magnetostrictive Actuators for Cryogenic Applications
- 57 Noncavitating Pump for Liquid Helium

Books and Reports

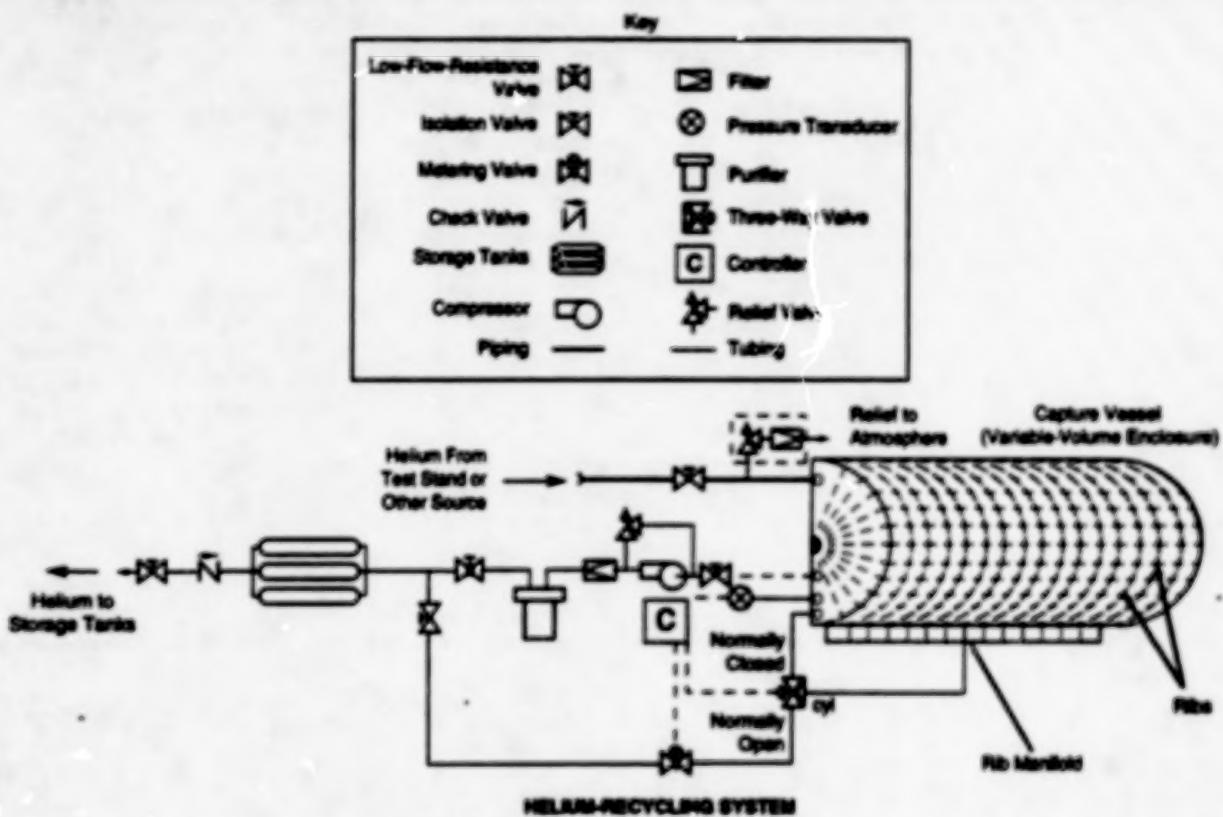
- 58 Experimental Study of Rotating Unbalanced-Mass Actuators

BLANK PAGE

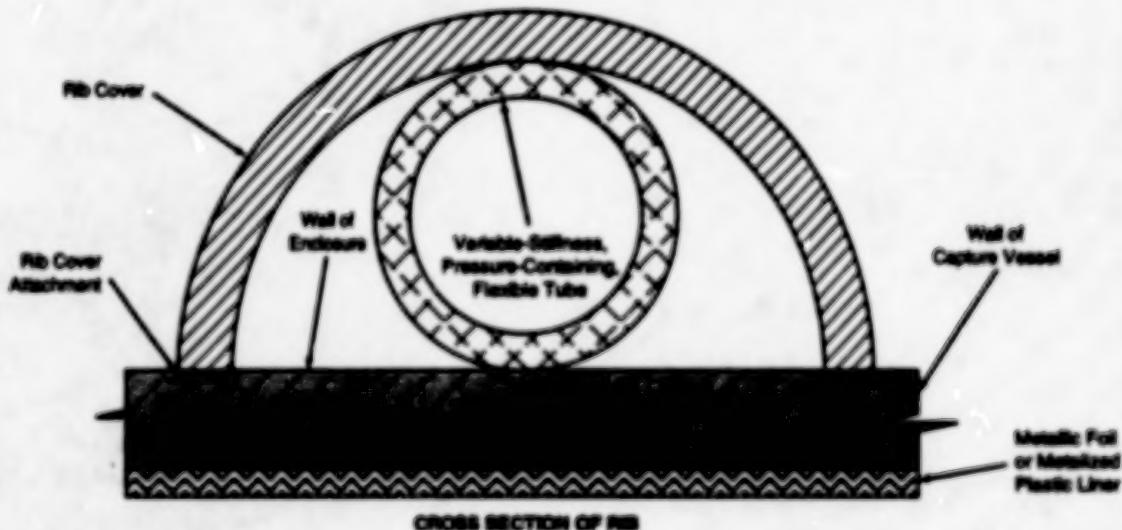
Helium-Recycling Plant

Helium would be conserved with little contamination or back pressure on the source.

Lyndon B. Johnson Space Center,
Houston, Texas



HELIUM-RECYCLING SYSTEM



Helium From a Test Stand would enter a variable-volume capture vessel for temporary storage before being transferred to tanks for longer-term storage. The capture vessel would be a buffer that would prevent the application of excessive back pressure to the test stand.

A proposed system would recover and store helium gas for reuse. The system would maintain the helium at 99.99-percent purity, preventing water vapor from the atmosphere or lubricating oil from pumps from contaminating the gas. The system would take in the gas at a nearly constant low back pressure near atmos-

pheric pressure [14.7 psi (0.101 MPa)]; that is, it would introduce little or no back pressure into the source of helium.

The helium-recycling system was conceived for test facilities of NASA and its contractors, which together use almost 10 million standard cubic feet (about 283,000 standard cubic meters) of helium

per year. Recycling would save an estimated \$650,000 (in 1992 prices) annually. The concept could also be extended to recycling of other gases. The helium to be recycled would enter the system through inlet pipes that would route the helium from a test stand or other source to a semicylindrical capture vessel (see figure).

The pipes, typically 3 in. (7.6 cm) in diameter, would be designed to offer low resistance to flow and, consequently, low pressure drop. The inlet plumbing would include a butterfly-type isolation valve and a filtering relief valve, which would keep out particulate contamination while preventing excessive buildup of pressure in the vessel.

The capture vessel would be a lined variable-volume enclosure that would be made to expand to accommodate the incoming helium. The growth of the capture vessel during intake of helium would prevent the buildup of back pressure that would interfere with the operation of the source of helium. The maximum expansion of the vessel would accommodate 60,000 ft³ (about 1,700 m³) of helium — the amount emitted by a full day's operation of the test stand for which the system was conceived.

The walls of the capture vessel would be made of a low-outgassing material

that would have a low permeability by helium, so as to minimize contamination and leakage of helium. A liner of metallic foil or metalized plastic would reduce contamination and leakage further. The bottom of the vessel would be tied down to prevent wind from blowing it away.

A volume-control system would monitor the pressure in the capture vessel and adjust the volume of the vessel accordingly by regulating the stiffness of an integral rib structure. The controller could be an electronic computer or a pneumatic logic circuit. It would use some of the recycled helium to pressurize the ribs through a metering valve or pressure regulator and a rib manifold. Increasing the pressure in the ribs would stiffen them, raising the enclosure and thereby increasing the volume. Decreasing the pressure would soften the ribs, lowering the enclosure, and thereby reducing the volume.

A compressor would pump the helium out of the capture vessel overnight

into storage in tanks at a pressure above 2,000 psi (14 MPa). The compressor would be of a diaphragm type that minimizes contamination. On its way to the tanks, the helium would be purified further in a liquid-nitrogen cold trap, which would remove water vapor and other impurities. The helium could also be filtered through porous metal to remove particles larger than 100 micrometers.

This work was done by Joseph Cook of Johnson Space Center. Further information is contained in a TSP [see page 1].

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, Johnson Space Center [see page 1]. Refer to MSC-22091.

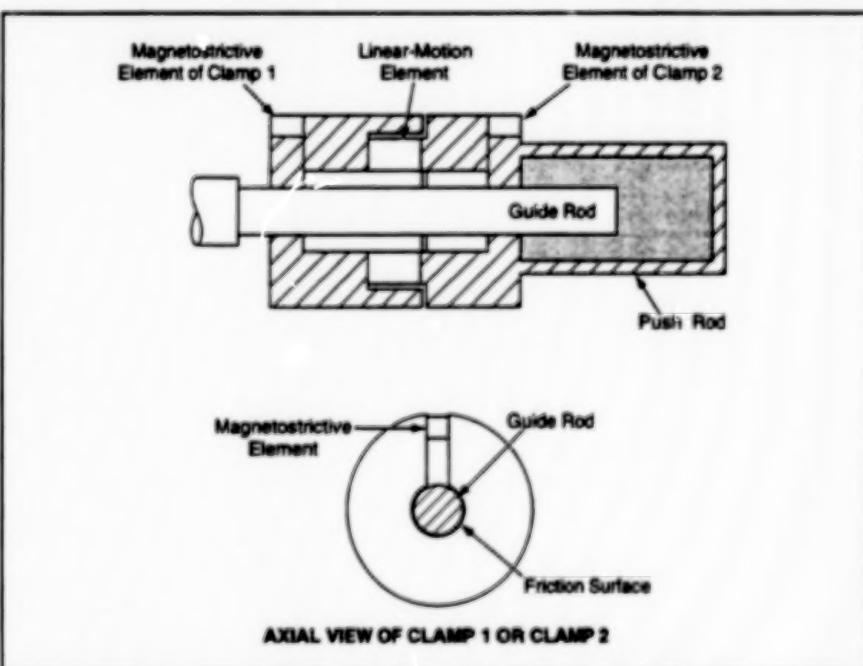
Magnetostrictive Actuators for Cryogenic Applications

These actuators would work at both room and cryogenic temperatures.

Linear-translation motors containing magnetostrictive actuator elements have been proposed for use in making fine position adjustments on scientific instruments at temperatures from near absolute zero to room temperature. Like piezoelectric actuators of a similar type described previously in *NASA Tech Briefs*, these actuators would produce small increments of linear motion and would operate in a "set-and-forget" mode in the sense that they would automatically lock themselves against motion when power was not applied. Conversely, they would not consume or dissipate power when stationary.

The reason for using magnetostrictive (as distinguished from piezoelectric) actuators is that magnetostrictive actuators function throughout the desired temperature range, whereas piezoelectric actuators tend to become inoperable in cryogenic environments. Although voice-coil actuators have been demonstrated at low temperatures, they cause waste of cryogen (because of continued consumption and dissipation of power when stationary at nonequilibrium positions), offer limited maximum displacements, and require complicated controllers. Because of the "set-and-forget" mode, the proposed actuators would thus offer advantages over voice-coil actuators; namely, less power dissipation (and consequently less

NASA's Jet Propulsion Laboratory, Pasadena, California



AXIAL VIEW OF CLAMP 1 OR CLAMP 2

When this Linear-Translation Motor was not powered, friction created between the clamps and the guide rod would maintain the position of the motor along the guide rod. The motor would move in steps; during each step, a motor controller would systematically turn the power on and off in the magnetostrictive elements.

consumption of cryogen). The proposed linear-translation motors could also be made to produce large maximum displacements.

Moreover, the controller of a motor of the proposed type could be relatively sim-

ple because it would be required to generate only simple sequences of on/off commands for three magnetostrictive actuator elements. As shown in the figure, these elements would be (1) a magnetostrictive element that, when energized, would pry

clamp 1 open just enough to allow clamp 1 to slide on the guide rod; (2) a similar element in clamp 2; and (3) a linear-motion element.

One cycle of operation (during which the motor would effect one increment of linear motion) would consist of the following events:

1. Friction in the clamps would maintain the initial position as long as power was not applied.
2. Power would be applied to the magnetostrictive element in clamp 1, causing this clamp to release the guide rod while clamp 2 continued to grip.
3. The linear-motion element would be energized, causing clamp 1 to translate one increment of position, relative to clamp 2, along the guide rod.
4. Power would be removed from the magnetostrictive element in clamp 1, causing clamp 1 to grip the rod. Thus, at this point, both clamps would grip the rod.
5. Power would be applied to the magnetostrictive element in clamp 2, causing this clamp to release the guide rod while clamp 1 continued to grip.
6. Power would be removed from the linear-motion element, causing clamp 2 to move along the guide rod by the same amount and in the same direction as clamp 1 moved previously.
7. Power would be removed from clamp 2, so that both clamps would grip the rod.
8. At this point, the motor would be unpowered at the incremented position, and friction between the clamps and the guide rod would maintain the position as long as power was not applied.

This work was done by Benjamin P. Dolgin of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1]. NPO-19218

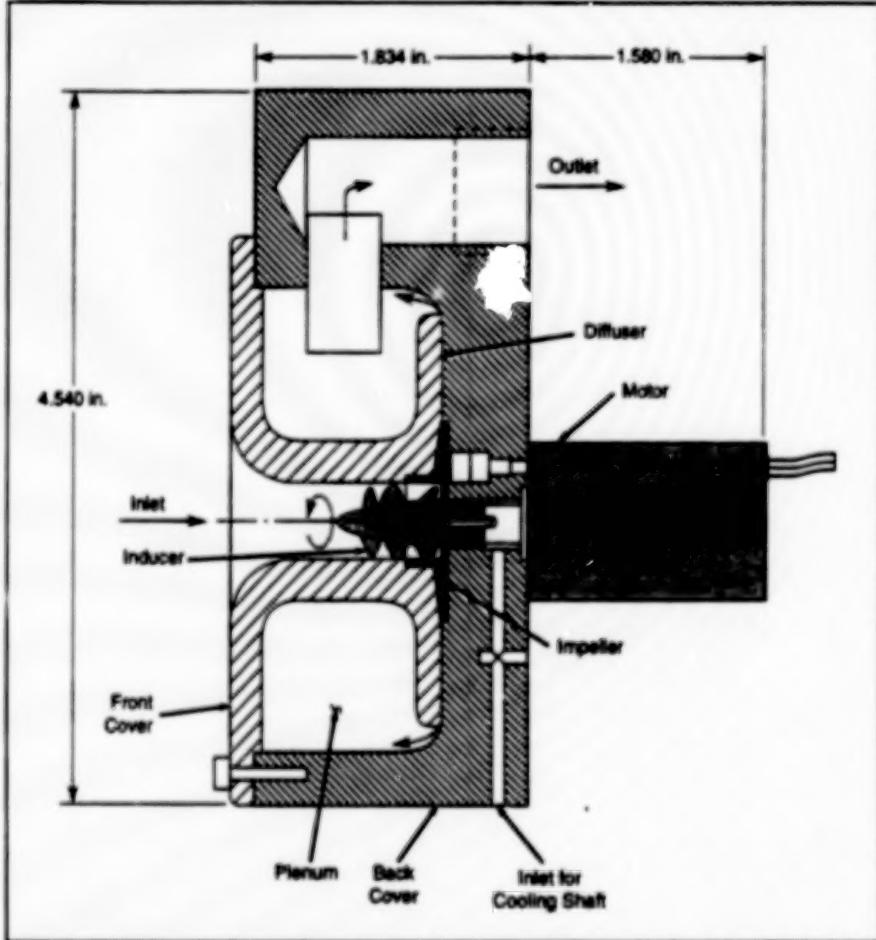
Noncavitating Pump for Liquid Helium

An immersion pump features high efficiency in cryogenic service.

A simple and reliable centrifugal pump transfers liquid helium with a mass-transfer efficiency of 99 percent — no more than 1 percent of the cryogenic liquid is lost by boiling off due to energy input from the pump. The pump operates without cavitation even with negative values of net positive suction head. (Net positive suction head is the difference in static head between the liquid at the inlet of the pump and the static head in saturated vapor.)

The purpose of the pump is to transfer liquid helium from a supply to a receiver vessel, or to provide liquid helium flow for testing and experimentation. The pump includes an induction motor, a centrifugal impeller, a helical inducer, solid-lubricated ball bearings, and a housing (see figure). It is located in a supply vessel submerged in liquid helium, which also serves to cool the motor. With a net positive suction head of -12 cm or more, the pump delivers 800 L/h at a head of 128 J/kg and an impeller speed of 12,000 r/min.

Liquid helium first flows into the inducer, then flows into the centrifugal impeller. The inducer resists blockage of flow from cavitation and imparts a small head to the inlet stream to reduce the tendency to cavitation in the impeller. The centrifugal force in the impeller provides most of the pressure rise through the pump. The liquid helium flows radially away from the impeller through the diffuser region between front and back covers, then enters a plenum that surrounds the inlet. As the helium decelerates in the diffuser and plenum, its kinetic energy is partially converted into increased pressure. From the plenum, the liquid leaves the pump through the outlet port.



Liquid helium is drawn into the pump by the helical inducer, the primary function of which is to pressurize the helium slightly to prevent cavitation when the liquid enters the impeller. The impeller then pressurizes the liquid.

There are two ball bearings, that carry axial and radial shaft loads. The bearings consist of stainless-steel balls and raceways with fiberglass/polytetrafluoroethylene/nylon composite retainers, selected for a combination of durability and lubri-

cating properties in a severe cryogenic environment. After 24 h of continuous operation, the bearings showed light wear but no sign of imminent failure.

The motor is of a three-phase ac, squirrel-cage type. An induction motor

was chosen for its simplicity and absence of commutating electrical connections (graphite brushes in commutators are short-lived in a dry environment with no moisture for lubrication).

This work was done by Robert Hasenbain, Michael Izenson, Walter Swift, and Herbert Sixsmith of Creare, Inc., for Ames Research Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Ames Research Center [see page 1]. Refer to ARC-13229.

Books and Reports

Experimental Study of Rotating Unbalanced-Mass Actuators

Results confirm the expectation that scanning by use of RUM actuators can save power.

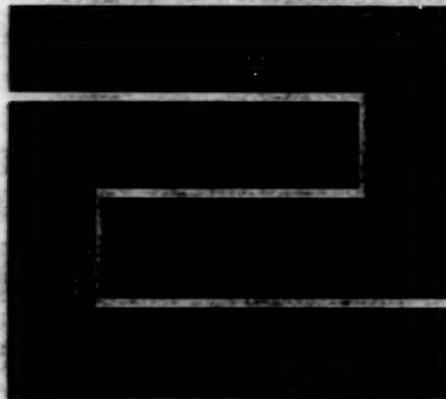
A report discusses the theory of rotating unbalanced-mass (RUM) actuators and describes experiments to test the concept of using RUM actuators for mechanical scanning of scientific instruments in linear, circular, or raster patterns. This concept and some aspects of the design and operation of RUM actuators have been described in several previous articles in

NASA Tech Briefs. In the specific RUM configuration that was investigated in these experiments, pairs of RUM actuators were used to generate basic tilt-scanning motions of a gimbaled beam, while gimbal torque motors were used both to keep the scans centered and as alternative scanning actuators for comparison. Previous calculations had shown that scanning by use of RUM actuators should consume significantly less power than does scanning by use of gimbal actuators in both gravitational and nongravitational environments. The results of these experiments were found to be in fair agreement with data from simplified computer simulations of the experiments. The results verify expectation of reduced power consumption; depending on the particular test configuration, the

power needed to scan with gimbal actuators ranged from 17 to 494 times that needed to scan with RUM actuators.

This work was done by Dean C. Alhom and Michael E. Polites of Marshall Space Flight Center. To obtain a copy of the report, "Results of a Laboratory Experiment That Tests Rotating Unbalanced-Mass Devices for Scanning Gimbaled Payloads and Free-Flying Spacecraft," see TSP's [page 1].

This invention is owned by NASA, and a patent application has been filed. Inquiries concerning nonexclusive or exclusive license for its commercial development should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-28994.



Fabrication Technology

Hardware, Techniques, and Processes

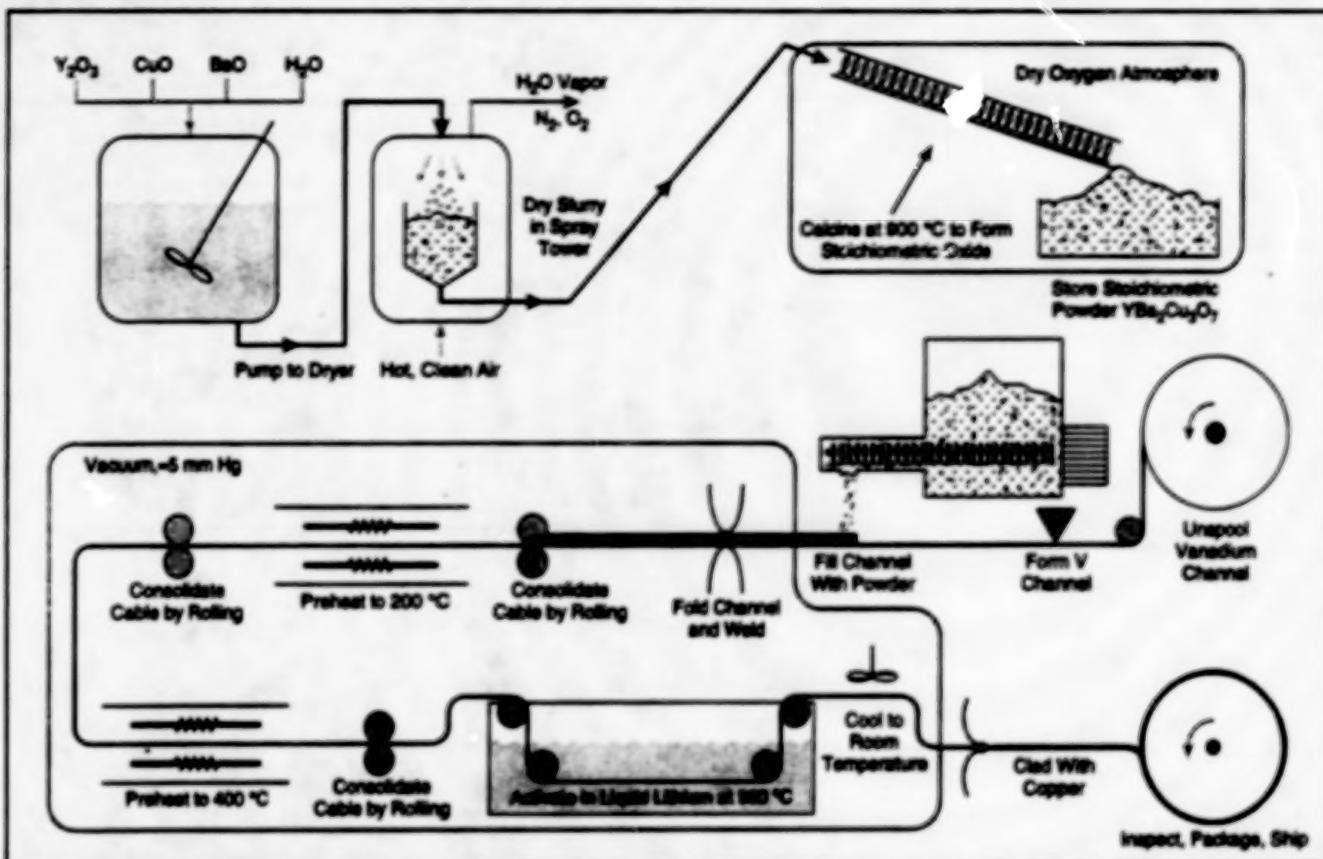
- 61 Manufacturing Superconducting Cables
- 62 Compressor Case Made With Filaments Wound With V-CAP Resin
- 62 Electrical-Discharge Machining of Perpendicular Passages
- 63 Brazing Dissimilar Metals
- 63 Verifying Removal of Red Penetrant Dye From Inspected Welds
- 64 Mathematical Model of Variable-Polarity Plasma Arc Welding

BLANK PAGE

Manufacturing Superconducting Cables

For protection, superconducting ceramic would be encapsulated before activation.

NASA's Jet Propulsion Laboratory,
Pasadena, California



A Cable Containing Superconducting Ceramic sheathed in metal would be drawn in a continuous process. The ceramic would be converted to a high-temperature superconductor after it had been encapsulated in metal foil.

A process has been proposed for the manufacture of cables that contain the ceramic high-temperature-superconductor YBa₂Cu₃O_{7- α} . The cables would carry electrical current with little or no loss of power when cooled to or below temperatures of about -200 °C. The process would accommodate the brittle nature of YBa₂Cu₃O_{7- α} and would be economical and readily controllable. It would also be flexible in the sense that it could be modified to accommodate a variety of precursor materials to be processed into YBa₂Cu₃O_{7- α} .

The process (see figure) would begin with mixing of oxides of yttrium, barium, and copper with water to form a slurry, followed by drying of the slurry, followed by calcining in a rotating kiln to obtain a stoichiometric precursor powder oxide, YBa₂Cu₃O₇. The precursor powder would then be cooled, stored, and inspected for quality. A dry oxygen atmosphere would be provided during calcining, cooling, storage, and inspection to protect the precursor powder from water and other atmospheric contaminants.

In the next few steps, a thin metal foil would be unspooled and formed into a V-channel, into which the precursor powder would be metered. The foil is required to be permeable by oxygen at the high temperature (850 to 950 °C) to which the cable would subsequently be exposed (to activate the superconductivity as explained below), but to be impervious to oxygen and other environmental contaminants at lower temperatures at which the cable would be used. Vanadium has been tentatively selected as the foil material because it exhibits the desired permeability-versus-temperature characteristic, plus a favorable thermal-expansion characteristic. Niobium and tantalum may be suitable alternatives as foil materials. The foil would be continuously folded over and welded to seal the powder inside. The filling, folding, sealing, and most of the following operations would be performed in a vacuum to prevent trapping of gaseous oxygen in the cable. The cable would then undergo several steps of consolidation by rolling and heating, and the heating in these

steps would also constitute preheating for the following activation step.

Activation of superconductivity would involve a controlled combination of (1) heating the precursor stoichiometric YBa₂Cu₃O₇ powder to the activation temperature and (2) removing a small amount of oxygen to obtain the desired superconducting ceramic that is slightly substoichiometric in oxygen; namely, YBa₂Cu₃O_{7- α} (where α represents the deviation from stoichiometry). For this purpose, the cable would be pulled through a bath of molten lithium at the upper activation temperature of 950 °C. The affinity of lithium for oxygen would create a thermodynamic gradient of oxygen activity across the encapsulating foil, causing oxygen to diffuse out of the cable. Liquid sodium or liquid potassium might be used instead of liquid lithium, but their vapor pressures are higher — a disadvantage in the vacuum environment. The speed of the cable through the activation bath could be adjusted to adjust the immersion time and thus the amount of oxygen removed.

After emerging from the activation bath, the activated cable would be slowly cooled to room temperature, then brought out of the vacuum, then optionally clad with copper, which would provide

additional support and protection and would serve, along with the foil, as an alternative path for current if the superconductor should fail. Finally, the cable would be wound on a spool for shipment.

This work was done by Christopher England of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1]. NPO-19196

Compressor Case Made With Filaments Wound With V-CAP Resin

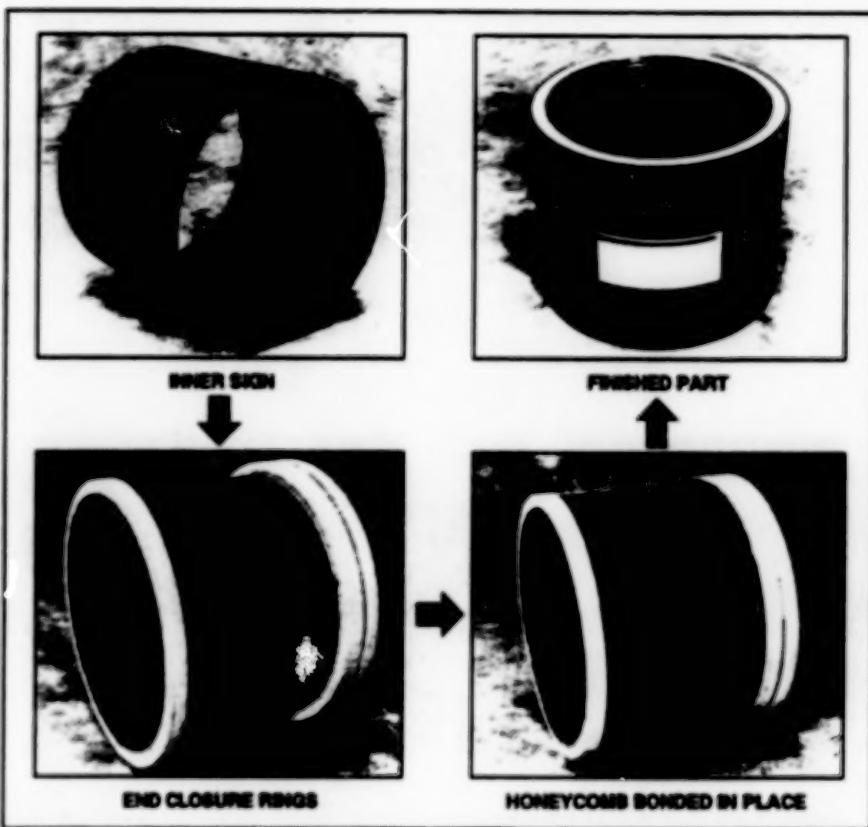
Weights and costs could be reduced.

Lewis Research Center, Cleveland, Ohio

A laminated cylindrical composite-material compressor case has been fabricated in a process that includes winding of filaments wetted with matrix resins. This case is a prototype of lightweight composite compressor cases that are required to withstand internal temperatures as high as 700 °F (371 °C). These composite cases are candidates for replacing titanium compressor cases in high-temperature turbines: it has been estimated that in comparison with titanium compressor cases, the composite compressor cases could weigh about 30 percent less and cost about 10 percent less.

The figure shows the prototype case as a workpiece at various stages of fabrication. The case includes a two-layer inner skin, which was formed by winding filaments wetted with VCAP-75 polyimide resin onto a cylindrical mandrel. The filaments in the innermost layer were made of glass and were wound in a hoop configuration; the filaments in the adjacent layer were made of graphite and were wound in three sublayers in hoop, helical, and hoop configurations, respectively. The layers and sublayers were cured by heating, along with a combination of externally applied pressure and internal vacuum to remove bubbles.

The cured inner skin was trimmed to length and removed from the mandrel, at which point it was in the form of a thin-walled, self-supporting cylinder. Its outer surface was grit-blasted to roughen it to enhance adhesion of the following layers. Premolded closure rings made of phenolic resin, chopped glass, and microballoon filler were installed on the ends of the cylindrical inner skin and bonded in place by use of a high-temperature adhesive in bonding grooves. Next, a film adhesive



The Composite Compressor Case is fabricated in stages that include winding of filaments wetted with matrix resins, adhesive bonding, layup, and several intermediate curing steps.

was installed on all bonding surfaces and a premachined aromatic polyamide honeycomb layer was laid up onto the cylindrical inner skin. The honeycomb core was wrapped with shrink tape and the film adhesive was then cured.

The workpiece was reinstalled on the mandrel, more film adhesive was applied, and an outer layer of graphite filaments wetted with a high-temperature epoxy was wound onto the outer surface of the honeycomb core in sublayers with hoop

and helical configurations. The epoxy was then cured, then the workpiece was trimmed to length and removed from the mandrel in its final form.

This work was done by Raymond Vannucci and James Sutter of Lewis Research Center and W. Donald Humphrey, A. John Ayorinde, and Jeremy Eaton of Brunswick Composites, Inc. No further documentation is available.

LEW-15990

Electrical-Discharge Machining of Perpendicular Passages

A perpendicular telescoping electrode reaches into previously inaccessible places.

A perpendicular telescoping electrode is used to perform electrical-discharge machining (EDM) of an internal passage

through the previously inaccessible depth of a metal workpiece. More specifically, it is used to make an internal passage per-

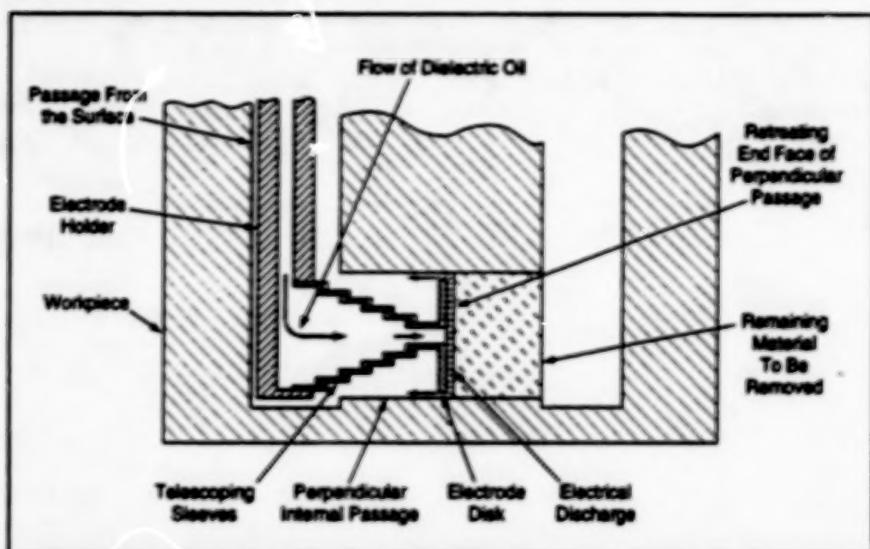
Marshall Space Flight Center, Alabama

pendicular to a passage that enters from the outer surface. Heretofore, the only way to make such a perpendicular inter-

nal passage was to make an access hole, then machine through the depth of material to form the passage, then restore the overlying material by welding a plug or other appropriately dimensioned piece of metal over the access hole.

The telescoping electrode is mounted on the lower end of an EDM electrode holder that is lowered into the passage from the surface (see figure). The telescoping electrode consists of an electrode disk supported on the narrowest one of a series of progressively narrower thin-walled copper sleeves. Attached to the smallest sleeve is an electrode disk the same diameter as the largest sleeve. The disk is the actual burning area. The sleeves are assembled concentrically and are compressed axially before lowering the electrode into the passage from the surface.

The telescoping electrode is positioned at the depth of the perpendicular passage to be machined and its axis is aligned with the axis of the perpendicular passage. The electrode holder is hollow and is used as a manifold to supply dielectric oil to the electrode. The pressure exerted by the oil flowing through the electrode pushes the sleeves out-



The Telescoping Electrode is extended into the perpendicular passage by the pressure exerted by the flowing dielectric oil.

ward, thus extending the electrode so that the burning surface of the electrode disk follows the retreating end face of the passage. After the EDM is complete, the electrode can be retracted by use of a vacuum or compressed air.

This work was done by R. Michael Malinak and Gary N. Booth of Rockwell

International Corp. for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-30027.

Brazing Dissimilar Metals

Metals with widely different coefficients of thermal expansion can be joined.

Brazing has been found to be an effective technique for joining ordinary structural metals to brittle, low-thermal-expansion refractory metals. Specifically, a brazing process has been established for (a) joining copper or nickel flanges to the ends of vacuum-plasma-sprayed tungsten tubes and for (b) joining stainless-steel flanges to the ends of tubes made of an alloy of molybdenum with 40 percent of rhenium.

In the original application, the tubes were furnace-cartridge tubes 23.5 in. (59.7 cm) long and 1 in. (25.4 mm) in diameter, with wall thickness of 0.030 in. (0.76 mm). The tubes were of test-tube configuration and the flanges were joined to the open ends.

In preparation for brazing, the tubes were cleaned of surface contaminants by placing them in a hydrogen furnace in the temperature range of 2,000 to 2,100 °F (approximately 1,100 to 1,150 °C) for 15 minutes. That area on each tube that was to be joined to a flange was then pre-tinned by wrapping the area with a thin brazing foil and heating it to the melting temperature.

After pretinning, each flange was machined to provide a braze-joint gap of 0.003 to 0.004 in. (0.076 to 0.10 mm). Each tube was then put together with its flange, and a small amount of brazing alloy was placed at the opening of the braze joint to provide sufficient material to fill the gap. A braze-stop-off solution was applied to the

Marshall Space Flight Center, Alabama

surfaces surrounding the joint to prevent the brazing alloy from flowing onto these surfaces and thereby make it unnecessary to perform a postbrazing cleanup. The tubes were then brazed, variously, in an inert or vacuum environment in the vertical orientation with the flanges at the lower ends.

This work was done by Philip D. Krotz, William M. Davis, and Daniel L. Wisner of Rockwell International Corp. for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Marshall Space Flight Center [see page 1]. Refer to MFS-30105.

Verifying Removal of Red Penetrant Dye From Inspected Welds

A clean surface is assured for more sensitive inspection with fluorescent penetrant dye.

A simple procedure has been devised to ensure that visible (red) penetrant dye that has been used to identify flaws in a welded surface has been completely removed from the surface. It is necessary to ensure

complete removal because any residual red dye could interfere with a subsequent inspection in which ultraviolet illumination and a fluorescent penetrant dye are used to identify smaller defects.

Marshall Space Flight Center, Alabama

The procedure consists in applying a reversible penetrant developer to the surface to be inspected. The developer contains a fluorescent dye that reacts with, and thus is eliminated by, the red penetrant

dye. Therefore, when the surface is viewed under ultraviolet illumination, the residual spots of red penetrant dye stand out as dark spots against a fluorescent background. Once the removal of red penetrant

dye has been completed and verified by this procedure, the developer is removed from the surface by rinsing with water.

This work was done by Jan R. Torkelson of Rockwell International Corp.

for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].
MFS-30001

Mathematical Model of Variable-Polarity Plasma Arc Welding

A mathematical model of the variable-polarity plasma arc (VPPA) welding process has been developed for use in predicting the characteristics of welds and thus serving as a guide for the selection of process parameters. These parameters include the welding electric currents in, and the durations of, straight and reverse polarities; the rates of flow of plasma and shielding gases; and the sizes and relative positions of the welding electrode, the welding orifice, and the workpiece. The model includes submodels that approxi-

mate the time-dependent major electrical, thermal, radiative, and mechanical phenomena involved in the transfer of energy through the electrode and gases to the workpiece and environment. The model predicts (1) electric potentials (in both straight and reversed polarities) at key locations; (2) electric power input at the electrode, within the orifice, in the free plasma jet column, and in the weld key-hole in the workpiece; (3) power loss at the electrode, within the orifice but outside the electrode, in the standoff column, and in

the workpiece; (4) enthalpy of the plasma arc; and (5) widths and heights of the crown and root of the weld bead. In a study, predictions by this model were found to agree fairly well with data on both normal and abnormal experimental welds.

This work was done by R. J. Hung of the University of Alabama in Huntsville for Marshall Space Flight Center. Further information is contained in a TSP [see page 1].
MFS-26311



Mathematics and Information Sciences

Hardware, Techniques, and Processes

- 67 Modified Mean-Pyramid Coding Scheme
- 68 Algorithm for a Self-Growing Neural Network
- 68 Fast Parallel Computation of Multibody Dynamics
- 69 Software for Multivariate Bayesian Classification

BLANK PAGE

Modified Mean-Pyramid Coding Scheme

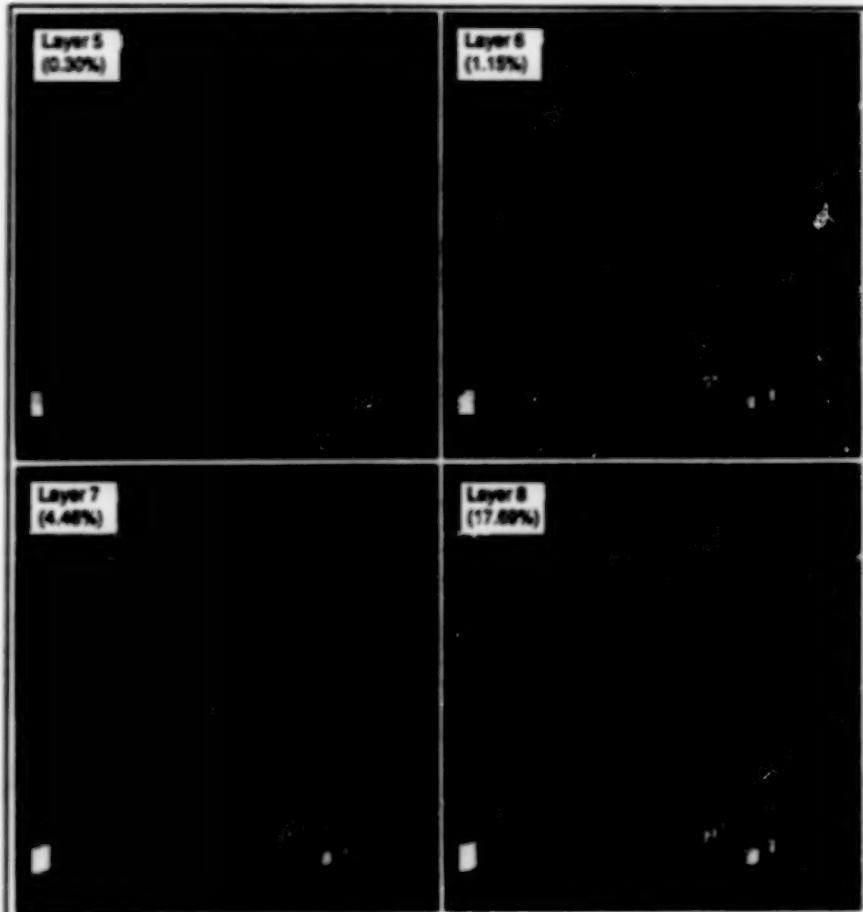
The data-expansion factor can be reduced from 1/3 to 1/12.

NASA's Jet Propulsion Laboratory,
Pasadena, California

A modified mean-pyramid coding scheme requires the transmission of slightly fewer data than does the unmodified mean-pyramid coding scheme. These are schemes for the progressive transmission of image data. In progressive transmission (which should not be confused with "progressive scanning"), data on the entire image area are transmitted in a sequence of frames in such a way that a coarse version of the image can be reconstructed after the receipt of the first frame and an increasingly refined version of the image can be reconstructed after the receipt of each subsequent frame.

Progressive transmission offers an advantage when the time and/or the bandwidth available for transmission of image data is limited. When using a traditional line-by-line or raster scan, it is usually necessary to transmit the entire image before the user at the receiving end can recognize it. In progressive transmission, the user can view the increasingly refined approximate received images and make a decision as to whether it is worthwhile to continue refining the image or to devote the limited communication channel to another purpose — possibly the transmission of another image. In the mean-pyramid coding schemes, progressive transmission is combined with data-compression coding to decrease the bandwidth of the data signal, thereby further increasing the efficiency of utilization of the communication channel.

In the unmodified mean-pyramid coding scheme, the image data are represented by a hierarchical structure, which, in its simplest form, involves the construction of each succeeding higher and coarser layer by averaging over blocks of pixels (typically 2×2) in the layer immediately below it. The highest layer (which represents the average pixel intensity of the entire image) is transmitted first, followed by the differences between the highest and second highest layers, followed in turn by the differences between the third and second layers, and so forth. The differences between adjacent layers tend to be smaller more often than they are larger and can, therefore, be encoded with fewer bits (that is, compressed). The re-



A **Recognizable Image** is formed after fewer than 5 percent of the image data have been transmitted, and an image of fairly high quality can be reconstructed with fewer than 20 percent of the data.

ceiver decodes the transmitted data and partly mimics the operation of the transmitter in reconstructing approximate images of successively finer resolution.

The unmodified mean-pyramid scheme partly counteracts the effort to compress data in that it increases the number of data to be transmitted for each layer. In the case of an 8-bit gray-scale image of 512×512 pixels, the increase is by a factor of 1/3. The modification is directed toward reducing this expansion, and it involves the reconstruction of the fourth pixel in each block from (a) the three other pixels in the block and (b) the mean pixel in the next higher layer. In the absence of any further correction, this technique of reconstruction could introduce a truncation error into the fourth pixel, inasmuch as the mean pixel is the truncated average of the 2×2 block.

The truncation error is either 0, 1, 2, or 3 levels and is quite evenly distrib-

uted. The needed correction is supplied by transmitting the error term, which can be coded with 2 bits instead of the full 8 bits. This correction and representation guarantees lossless transmission and reduces the data-expansion factor from 1/3 to 1/12. The modified mean-pyramid coding scheme also incorporates a context-based arithmetic code that compresses the difference and error terms. The figure illustrates the sequence of reconstructed images obtained by applying the modified mean-pyramid coding scheme to an oilfield.

This work was done by Kar-Ming Cheung and Richard Romer of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP [see page 1].
NPO-18890

Algorithm for a Self-Growing Neural Network

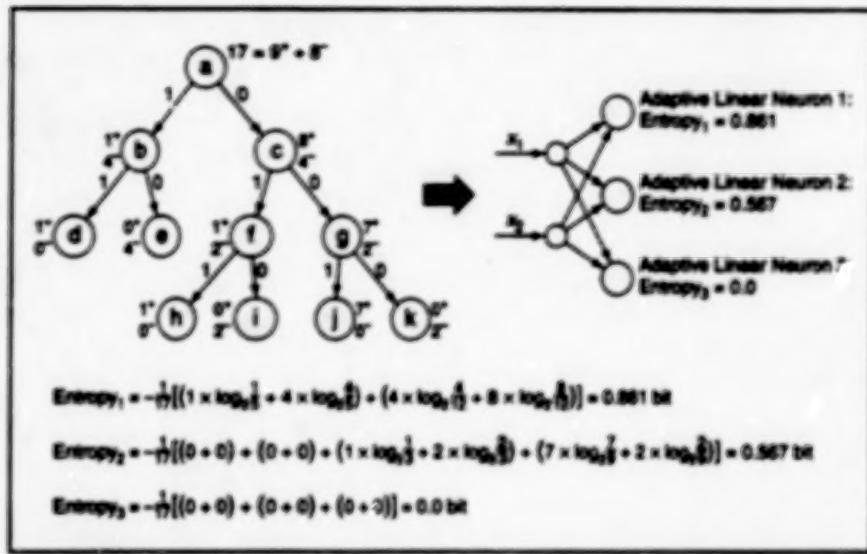
Either a crisp or a fuzzy measure on entropy can be used.

Lewis Research Center, Cleveland, Ohio

The CID3 algorithm simulates a self-growing neural network. It constructs decision trees equivalent to the hidden layers of the neural network (see figure). The algorithm is so named ("CID3" is short for Continuous ID3) because it is based on the ID3 algorithm, which dynamically generates a decision tree while minimizing the entropy of information. [The entropy of a piece of information, x , is given by $-p(x)\log_2 p(x)$, where $p(x)$ denotes the probability or frequency of occurrence of x .] The CID3 algorithm generates a feedforward neural network by use of either a crisp or a fuzzy measure of entropy.

A description of ID3 is prerequisite to a description of CID3. ID3 generates decision rules from a set of training examples. Each example is represented by a list of features. In the training process, classes to which the input data belong must be known. The idea is to find the minimum number of original features that suffice to determine these class memberships. ID3 uses information theory to select features that maximize the information gain in the sense that they maximize the decrease in entropy.

Inasmuch as the number of new nodes added to a decision tree depends on the number of values that a selected feature can take on, the ID3 algorithm requires features to have discrete values. The decision tree thus generated is then described in terms of hierarchical decision rules that must be used in the order specified by the tree structure. The condition part of a decision rule consists of a number of feature tests linked by AND/OR logical operators. The disadvantage of a feature test is that correlations between features are ignored. ID3 considers only the degree to which each



A Three-Node Hidden Layer of a neural network is generated from a decision tree, in the CID3 algorithm, by use of mathematical constructs called "adaptive linear neurons."

individual feature is significant for classification of training examples.

In the formulation of CID3, ID3 was modified to operate on continuous data and to search for weight vectors that minimize an entropy measure. This measure can be based on either the crisp entropy (the conventional information entropy defined above) or a fuzzy entropy. The fuzzy entropy measure is a modified entropy, derived from fuzzy-logic theory, that accounts for degrees or probabilities of membership in various classes.

The CID3 algorithm comprises the following major steps:

1. For a given problem with N training examples, start with a random initial weight vector, W_0 .
2. Using a learning-rule equation, a description of which would greatly exceed the space available for this article, search for a hyperplane in

weight-vector space that minimizes the applicable crisp or fuzzy entropy function at a given level of the decision tree or layer of the neural network.

3. If the minimized entropy is not zero, but is smaller than the previous value, add a node to the current layer of the neural network and return to step 2. Otherwise, go to step 4.
4. If the hidden layer consists of more than one node, generate a new layer that utilizes inputs from both the original training data and the outputs from all previously generated layers, and go to step 2. If the hidden layer consists of only one node, then the problem is reduced to a linearly separable one; stop.

This work was done by Krzysztof J. Cios of the University of Toledo for Lewis Research Center. Further information is contained in a TSP [see page 1].
LEW-15931

Fast Parallel Computation of Multibody Dynamics

Processors and computation time are utilized efficiently.

The constraint-force algorithm is a fast, efficient, parallel-computation algorithm for solving the forward dynamics problem (given the control forces, compute the motion) of a multibody system like a robot arm or a vehicle. This is the first known algorithm to solve the problem in a minimum time proportional to $\log(N)$ by use of an optimal number of processors proportional to N , where N is the number

NASA's Jet Propulsion Laboratory, Pasadena, California

Version	Time (in ms)	Abs. Acc.	Speed
Serial	481	1.00	-
N-Parallel	87	5.53	5.53
2N-Parallel	66	6.97	1.26
3N-Parallel	55	8.75	1.25

Significant Speedups Were Achieved in a test computation by use of three different parallel versions of the constraint-force algorithm.

of dynamical degrees of freedom: in this sense, the constraint-force algorithm is both a time-optimal and a processor-optimal parallel-processing algorithm.

Fast algorithms and computers have been sought for use in teleoperated and other control systems. In many of these systems, there is a need for faster-than-real-time simulation of multibody dynamics. Fast serial algorithms with processing

times proportional to N have been developed, but even the fastest serial algorithm does not exhibit the requisite computational efficiency, and further improvements in the efficiencies of serial algorithms seem unlikely.

Heretofore, in the development of a typical parallel-processing algorithm, it has been customary to first develop an algorithm for fast serial computation, then attempt to increase the computational efficiency through parallelization. A different approach was followed in the case of the constraint-force algorithm: From the outset, it was developed for efficient parallel computation.

The complexity of the underlying mathematics precludes a detailed description of the algorithm in this article. From the perspective of physics, a given algorithm for multibody dynamics can be classified on the basis of, and can be considered to originate from, its force-decomposition strategy. From the perspective of mathematics, the algorithm can be classified on the basis of that factorization of the mass matrix of the dynamical system that cor-

responds to the specific force-decomposition strategy. The constraint-force algorithm is accordingly so named because it is based on a novel decomposition of interbody forces that leads to an explicit computation of constraint forces.

Researchers had previously often argued that because constraint forces are nonworking forces, one should avoid calculating them for the sake of computational efficiency. Although this argument is valid for serial computation, it does not necessarily apply to parallel computation. In this case, the force-decomposition strategy leads to a new factorization of the mass matrix and this factorization, in turn, results in the efficient constraint-force algorithm.

A salient feature of the constraint-force algorithm is that while it is less efficient than other algorithms for serial computation, it is strictly efficient for parallel computation. In addition to its theoretical significance, this algorithm is highly practical and can be efficiently implemented on a variety of commercially available computers that feature multiple-instruc-

tion/multiple-data parallel architectures. In a test, the constraint-force algorithm was successfully implemented on a 32-node hypercube computer in which each node contains an 80386 microprocessor, an 80387 arithmetic coprocessor, and 4 MB of random-access memory. The performances of several different versions of the algorithm (a serial version and parallel versions with three different levels of parallelism) were compared with respect to an $N = 8$ test case (see table). The results of the test confirmed expectations that the constraint-force algorithm exhibits a high degree of parallelism in computation with minimum communication and synchronization requirements.

This work was done by Amir Fijany of Caltech and Gregory L. Kwan and Nader Bagherzadeh of the University of California at Irvine for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).

NPO-19340

Software for Multivariate Bayesian Classification

This program can be embedded in another program or run by itself.

PHD is a general-purpose classifier computer program. Originally part of another program, PHD has been extracted and made available for wider use.

"Classifier" in this context denotes a program that accepts data in the form of vectors and assigns a class label to each vector. Typically, a classifier is trained to recognize a set of classes by (a) presenting it with training data in the form of a set of vectors to which the correct labels have been attached, then (b) running a "learning" procedure to characterize the classes. Thereafter, when presented with new, unlabeled vectors, the classifier is required to determine the class to which each vector most likely belongs, along with a measure of confidence in the correctness of the classification of each vector.

PHD was constructed according to the concept of Bayesian inference, a mathematically optimal family of algorithms for classification that has long been a fundamental technique for classification of patterns. PHD uses Bayesian methods to classify vectors of real numbers, based on a combination of statistical techniques

that include multivariate density estimation, Parzen density kernels, and the EM (Expectation Maximization) algorithm. By means of a simple graphical interface, the user can train the classifier to recognize two or more classes of data and then use it to identify new data.

The classifier is trained to create a statistical model, known as a "mixture density," for each of several classes of training data. During training, in order to model the training data for a given class, the program searches for a maximum-likelihood (best-fitting) density estimate constructed of one or more multivariate Gaussian kernels. Overfitting (failing to distinguish low-probability noise from higher probability data) is avoided by applying a minimum-description-length criterion to determine how many Gaussian kernels are justifiable in representing the training data. Whereas traditional pattern classification applies a discriminant function to distinguish each class from the others, PHD models each class separately. Thus, outliers that do not belong to any of the known classes can usually be spotted and labeled as "unidentified."

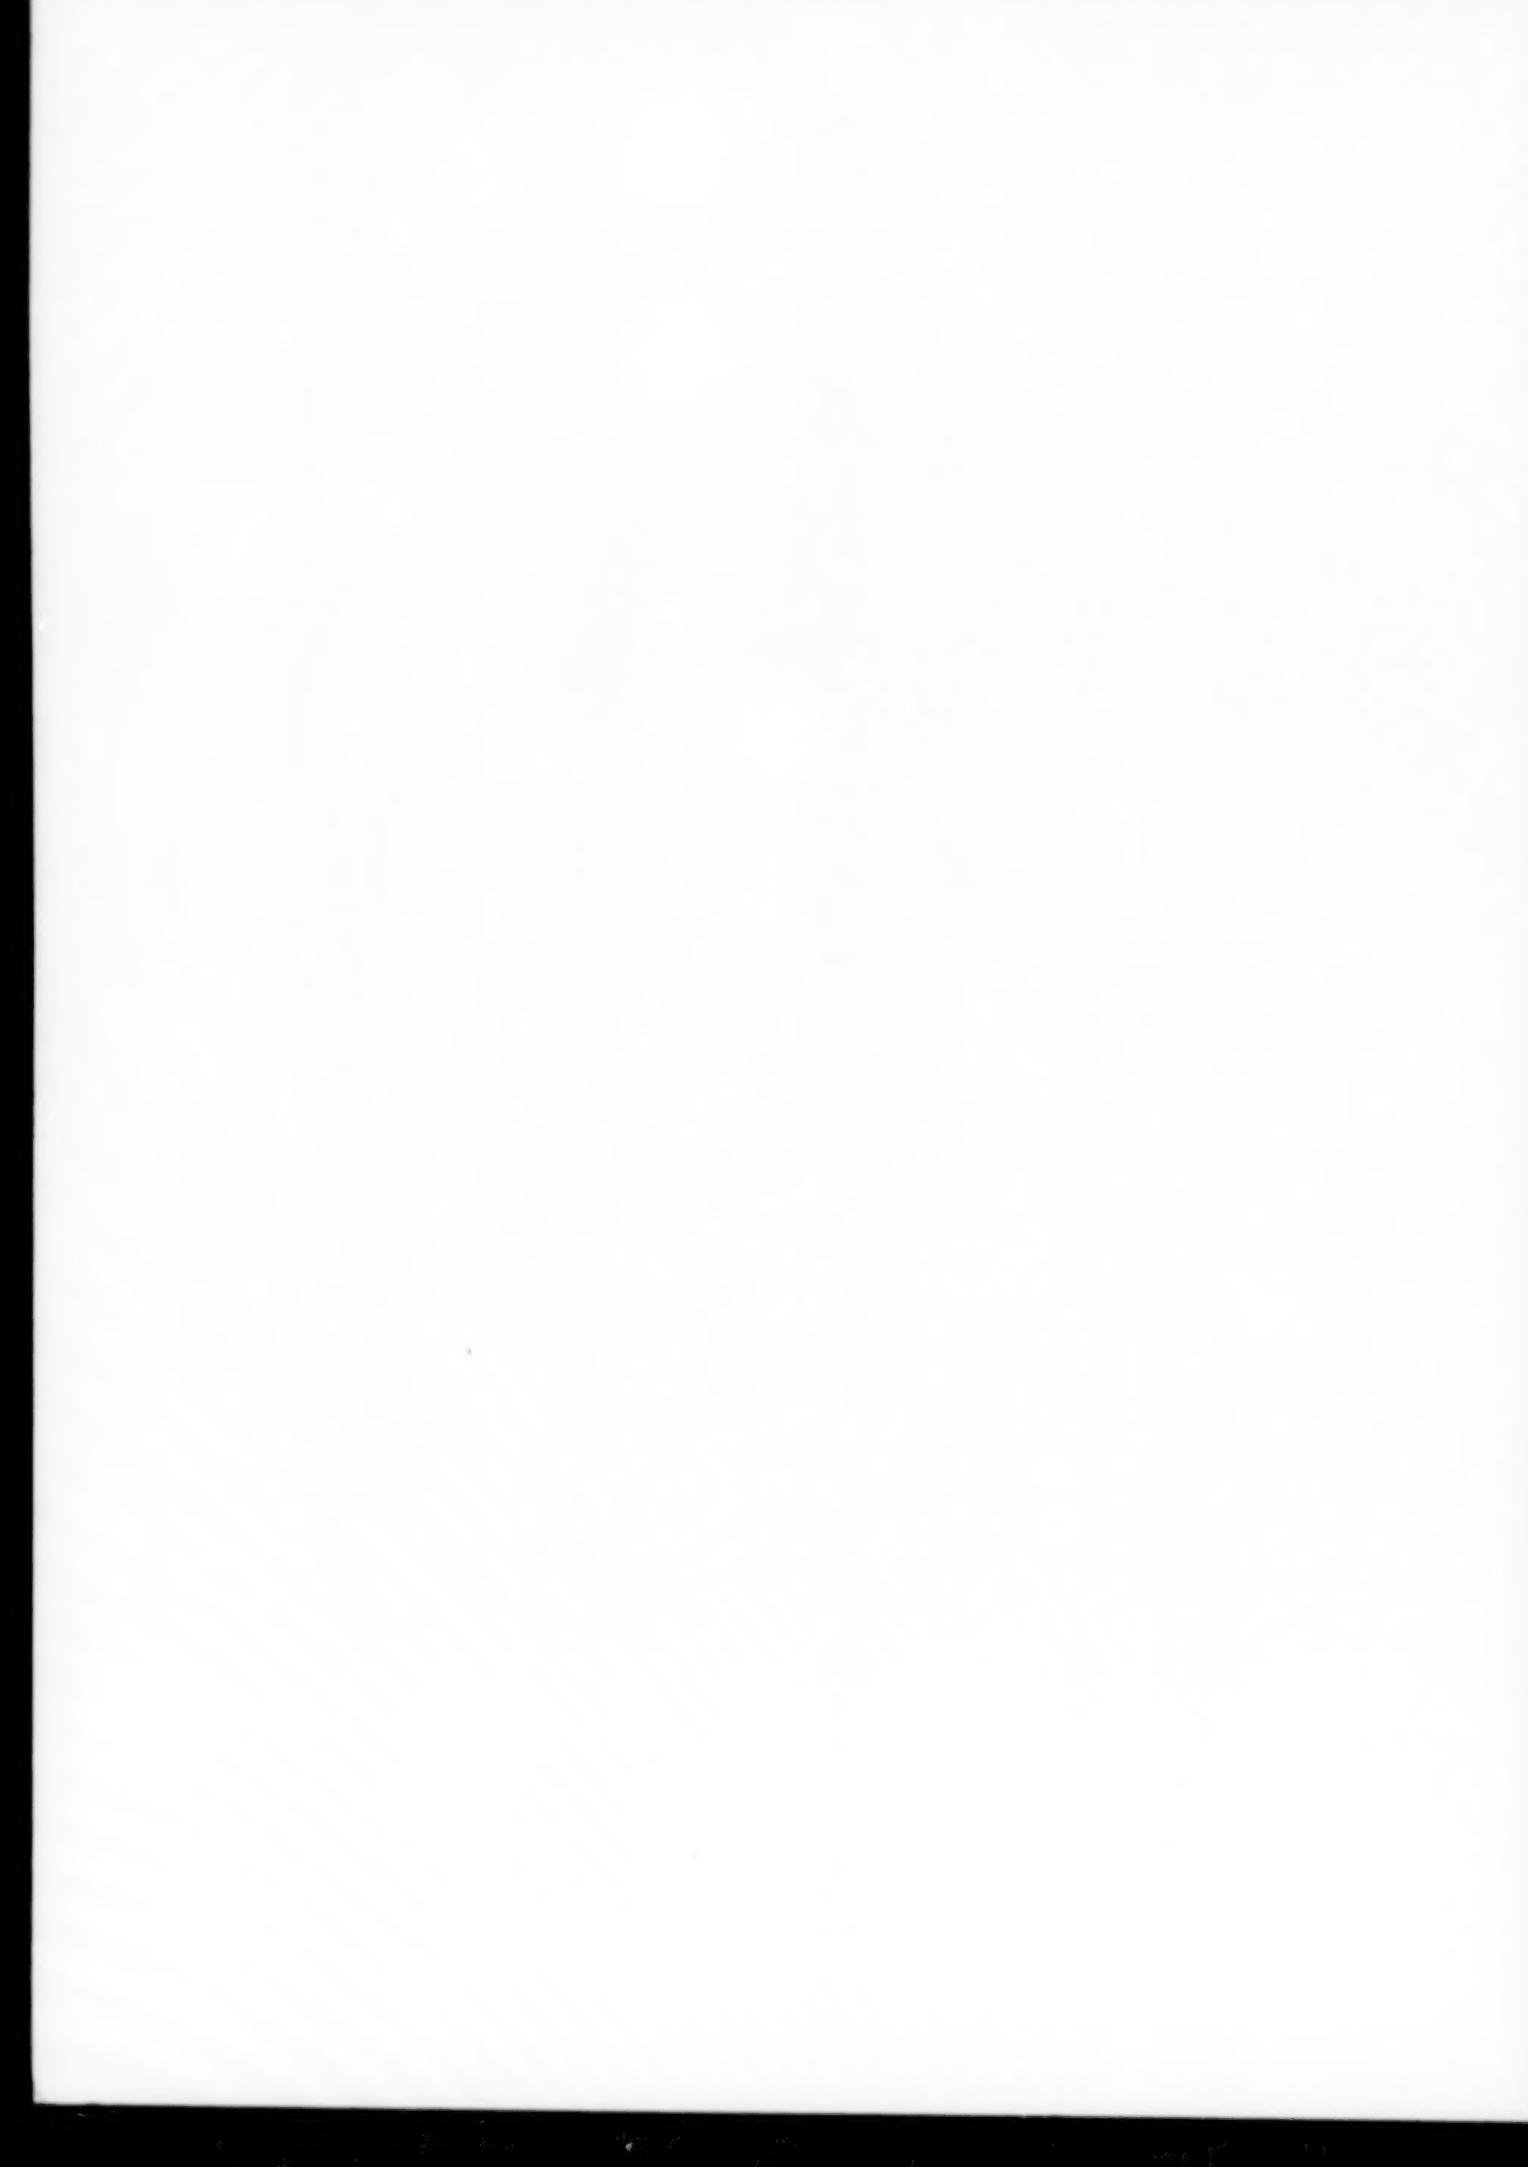
Ames Research Center,
Moffett Field, California

After finding densities for each of the training classes, PHD is ready to assign a class to an unlabeled vector by finding the highest likelihood for that vector from among the known class distributions. If the vector does not belong to any of the training classes, that fact will be reflected in the low confidence measure of its classification; thus, PHD is useful in a task that requires recognition of data that do not fit in any of the known classes.

PHD is written in ANSI C for Unix systems and is optimized for online classification applications. It can be embedded in another program, or it can run by itself using the simple graphical-user-interface. Online help files make the program easy to use.

This work was done by Ronald Saul and Philip Laird of Ames Research Center and Robert Shelton of Johnson Space Center. Further information is contained in a TSP (see page 1).

Inquiries concerning rights for the commercial use of this invention should be addressed to the Patent Counsel, Ames Research Center (see page 1). Refer to ARC-13383.



END

09.13.96